

Discrete Sampling and Interpolation: Universal Sampling Sets for Discrete Bandlimited Spaces

Brad Osgood *Member, IEEE* Aditya Siripuram William Wu
Information Systems Laboratory
Stanford University

Abstract—We study the problem of interpolating all values of a discrete signal f of length N when $d < N$ values are known, especially in the case when the Fourier transform of the signal is zero outside some prescribed index set \mathcal{J} ; these comprise the (generalized) bandlimited spaces $\mathbb{B}^{\mathcal{J}}$. The sampling pattern for f is specified by an index set \mathcal{I} , and is said to be a universal sampling set if samples in the locations \mathcal{I} can be used to interpolate signals from $\mathbb{B}^{\mathcal{J}}$ for any \mathcal{J} . When N is a prime power we give several characterizations of universal sampling sets, some structure theorems for such sets, an algorithm for their construction, and a formula that counts them. There are also natural applications to additive uncertainty principles.

Index Terms—Compressed sensing, Discrete Fourier transforms, Discrete time systems, Interpolation, Sampling methods, Uncertainty

I. INTRODUCTION

IN this paper and in a sequel [1] we consider the problem of interpolating all values of a discrete, periodic signal $f: \mathbb{Z}_N \rightarrow \mathbb{C}$, $N \geq 2$, when $d < N$ values of f are known. One solution is a discrete form of the classical Nyquist-Shannon theorem, where the spectrum of the signal is assumed to vanish outside a contiguous band of frequencies; see [2], for example. At the other extreme is the new and important area of compressed sensing, where no assumptions on the spectrum are made. For this, of the many papers we mention only [3], [4] and [5], since we will refer to this work later.

Our approach to the problem is in between, though we begin by formulating a very general definition.

Definition 1: Let \mathbb{Y} be a d -dimensional subspace of \mathbb{C}^N , let $\mathcal{I} \subset [0 : N - 1]$ be an index set of size d , and let $\mathcal{U}_{\mathcal{I}} = \{u_i : i \in \mathcal{I}\}$ be a set of d vectors in \mathbb{Y} . We say that $(\mathcal{I}, \mathcal{U}_{\mathcal{I}})$ is an *interpolating system* if each $f \in \mathbb{Y}$ can be written as

$$f = \sum_{i \in \mathcal{I}} f(i) u_i. \quad (1)$$

We call \mathcal{I} a *sampling set* and $\mathcal{U}_{\mathcal{I}}$ an *interpolating basis*. When we refer simply to a sampling set we always mean that it is

associated with an interpolating basis. If the vectors u_i are orthogonal we say that $(\mathcal{I}, \mathcal{U}_{\mathcal{I}})$ is an *orthogonal interpolating system* and that $\mathcal{U}_{\mathcal{I}}$ is an *orthogonal interpolating basis*.

The point of the definition is that the interpolation of all values of f uses the sampled values $f(i)$, $i \in \mathcal{I}$, which might be thought of as measurements of f with respect to the fixed, natural basis of the ambient space \mathbb{C}^N , while the basis $\mathcal{U}_{\mathcal{I}}$ is tailored to \mathbb{Y} and \mathcal{I} .¹ Note that \mathcal{I} need not consist of uniformly spaced indices, so the sampling may be irregular. Indeed, the results described here and in [1] were originally motivated by questions from colleagues in medical imaging who had observed that irregular sampling patterns could often give excellent results with less computation.

For us, to solve the interpolation problem for \mathbb{Y} is to find an interpolating system. It is a linear theory in all aspects. Every subspace has an interpolating system, though it may not be unique, but not every subspace has an orthogonal interpolating system. For a given subspace it is also not true that any index set is a sampling set for some interpolating basis, so the intervals between samples are not arbitrary. The only subspaces that have *orthonormal* interpolating systems are the coordinate subspaces. All of this is discussed in Section II. Orthogonal interpolating systems are the subject of [1], and we find interesting connections with difference sets, perfect graphs, tiling, and we answer affirmatively a discrete version of a conjecture of Fuglede.

In Section II we provide some basic results on interpolating systems in general. We quickly move, in Section III, to study bandlimited spaces, $\mathbb{B}^{\mathcal{J}}$, defined as signals whose discrete Fourier transforms are supported on \mathcal{J} . We do not require that \mathcal{J} be a set of contiguous indices, so this is more general than the situation in the discrete Nyquist-Shannon theorem (though we continue to use the term “bandlimited” for short).

In Section IV we begin to concentrate on universal sampling sets, namely index sets \mathcal{I} that are sampling sets for any bandlimited space $\mathbb{B}^{\mathcal{J}}$ with $|\mathcal{J}| = |\mathcal{I}|$. That is, \mathcal{I} is universal if the sampling pattern specified by \mathcal{I} can be used for interpolation of signals from any $\mathbb{B}^{\mathcal{J}}$. Universal sampling sets were used in [4] for multicoset sampling and in [5] in connection with compressed sensing. Here our central result gives several necessary and sufficient conditions for an index

Manuscript received August 31, 2011, revised January 20, 2012. A. Siripuram was supported by a Stanford Graduate Fellowship. W. Wu was supported by the Frank and Eva Buck Foundation

W. Wu is currently at the Jet Propulsion Laboratory, Pasadena, CA.

The authors are listed alphabetically.

Copyright (c) 2011 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

¹We could make the definition even more general and allow \mathbb{Y} to be a subspace of any finite-dimensional vector space \mathbb{X} , and sample $f \in \mathbb{Y}$ with respect to any fixed basis of \mathbb{X} , but the present definition suffices.

set to be universal when N is a prime power. A mathematical consequence of our result is a generalization of Chebotarev's theorem on the invertibility of submatrices of the Fourier matrix.

In Section V we show that a universal sampling set has an interesting structure as a disjoint union of what we call elementary universal sets, and through this analysis we are able to count the number of universal sampling sets of a given size. We also introduce maximal (and minimal) universal sampling sets which in turn enter naturally into the uncertainty principles that we discuss in Section VI. As an application of uncertainty and universality we prove a “random” uncertainty principle, and deduce a generalization of the Cauchy-Davenport theorem from additive number theory. Our debt to the work in [6] and [3] is clear. Many of our results assume that N is a prime power, and naturally we wonder whether this can be generalized.

The definitions we introduce and the methods we use are based primarily on properties of index sets when the elements are reduced modulo powers of a prime. With a few exceptions (e.g., minimal and cyclotomic polynomials) these can be considered elementary, and it is surprising (to us) how far they lead. The methods here also seem rather different from those of compressed sensing. In compressed sensing, which is nonlinear in theory and practice, the recovery of a signal from samples does not require knowledge of the frequency spectrum, whereas linear theories like ours cannot do without knowledge of the spectrum. Nevertheless, with universality the sampling patterns in our approach do not depend on the frequencies, the reconstruction of a signal from its samples is by linear operations, and the samples are “samples” in the classical sense instead of random projections of the signal onto a measurement basis as is done in compressed sensing. Both approaches start with discrete signals, but one needs to sample an analog signal in the first place and this analog sampling generally needs some knowledge of the frequency spectrum. Works such as [4] and [5] confront this issue through “spectrum blind” sampling, and they end up needing the idea of universality in the process. It is also interesting that the linear theory here can be used to prove a random uncertainty principle without the necessity of nonlinear techniques, though our result is not as strong as the result in [3]. We hope to pursue the connections and differences further. We refer to [2] and [7] for additional results, discussion, and examples. See also Appendix C for references to papers on universality for continuous-time signals.

II. GENERAL PROPERTIES, EXISTENCE OF INTERPOLATING SYSTEMS

This section is a summary of elementary properties of interpolating systems, including existence theorems in both an algebraic and geometric formulation. The ideas are simple enough, but they fit together nicely and are an essential foundation for the less simple work to follow.

We fix some notation. Without further comment we will identify a vector in \mathbb{C}^N with its N -periodic extension and vice versa, and we typically index vectors from 0 to $N - 1$.

(We assume periodicity because the discrete Fourier transform will soon enter the picture.) For $i \in [0 : N - 1]$ we let $\delta_i : \mathbb{Z}_N \rightarrow \mathbb{C}$ be the (periodized) discrete δ -function shifted to i , so that $\{\delta_0, \delta_1, \dots, \delta_{N-1}\}$ is the natural basis of \mathbb{C}^N . The components of a vector in \mathbb{C}^N will always be in terms of the natural basis, but any fixed basis of \mathbb{C}^N would do for the following development. If $\mathcal{I} \subset [0 : N - 1]$ we let

$$\mathbb{C}^{\mathcal{I}} = \text{span}\{\delta_i : i \in \mathcal{I}\}.$$

Our first goal is to establish

Theorem 1: Any subspace \mathbb{Y} of \mathbb{C}^N has an interpolating system.

We will give two proofs, one geometric and one algebraic, and both are straightforward.

In the following, \mathbb{Y} is always a subspace of dimension d and \mathcal{I} is always an index set of size d . Let $\mathcal{I}' = [0 : N - 1] \setminus \mathcal{I}$. We record several facts.

An interpolating basis for a subspace \mathbb{Y} is trying to be the natural basis in the slots specified by the index set. In fact this is a characterization of interpolating bases.

Proposition 1: (i) A basis $\mathcal{U} = \{u_i : i \in \mathcal{I}\}$ for \mathbb{Y} is an interpolating basis if and only if

$$u_j(i) = \delta_j(i) \quad i, j \in \mathcal{I}.$$

(ii) Any natural basis vector δ_k lying in \mathbb{Y} is an element of any interpolating basis of \mathbb{Y} .

(iii) An interpolating basis is determined by its index set, more precisely, if $\{u_i : i \in \mathcal{I}\}$ and $\{v_i : i \in \mathcal{I}\}$ are interpolating bases for \mathbb{Y} then $u_i = v_i$ for all $i \in \mathcal{I}$.

Expanding on the first point in Proposition 1, the elements of an interpolating basis are perturbations of the natural basis vectors by vectors outside \mathbb{Y} :

Proposition 2: (i) Any interpolating basis $\{u_i : i \in \mathcal{I}\}$ of \mathbb{Y} is of the form

$$u_i = \delta_i + v_i,$$

where $v_i \in \mathbb{C}^{\mathcal{I}'}$. If $v_i \in \mathbb{Y}$ then $v_i = 0$.

(ii) The subspaces of \mathbb{C}^N having an orthogonal interpolating system are of the form $\mathbb{Y} = \text{span}\{\delta_i + v_i : i \in \mathcal{I}\}$ where the nonzero v_i are orthogonal vectors in $\mathbb{C}^{\mathcal{I}'}$.

We omit the proofs of Propositions 1 and 2. Part (ii) of Proposition 2 can be applied in the negative to find examples of subspaces that do not have an orthogonal interpolating basis – this is a much larger topic – and it also follows from part (ii) that the only subspaces having an orthonormal interpolating basis are the coordinate subspaces. Both of these points were raised in the introduction.

Geometric Proof of Theorem 1: It is easy to see that there is an index set \mathcal{J} of size $N - d$ such that $\mathbb{C}^N = \mathbb{Y} \oplus \mathbb{C}^{\mathcal{J}}$. Let $P : \mathbb{Y} \oplus \mathbb{C}^{\mathcal{J}} \rightarrow \mathbb{Y}$ be the projection of \mathbb{C}^N onto \mathbb{Y} along $\mathbb{C}^{\mathcal{J}}$. If $f \in \mathbb{Y}$ then, on the one hand,

$$f = \sum_{i=1}^N f(i) \delta_i.$$

On the other hand, since $\mathbb{C}^{\mathcal{J}} = \ker P$ and $Pf = f$ we have

$$f = Pf = \sum_{i=1}^N f(i) P\delta_i = \sum_{i \notin \mathcal{J}} f(i) P\delta_i.$$

Thus the $u_i = P\delta_i$ form an interpolating basis of \mathbb{Y} indexed by $\mathcal{I} = [0 : N - 1] \setminus \mathcal{J}$. ■

We see from this why an interpolating basis need not be unique. The ambiguity in choosing an interpolating basis arises from the ambiguity in choosing a complement; if there is not a unique choice of the complement $\mathbb{C}^{\mathcal{J}}$ of \mathbb{Y} , and generally there is not, then there is not a unique interpolating basis for \mathbb{Y} . However, the existence of an interpolating basis *produces* a complement to \mathbb{Y} :

Proposition 3: Let $\mathcal{U} = \{u_i : i \in \mathcal{I}\}$ be an interpolating basis of \mathbb{Y} . Then $\mathbb{C}^N = \mathbb{Y} \oplus \mathbb{C}^{\mathcal{I}'}$.

Proof: If we show that $\mathbb{Y} \cap \mathbb{C}^{\mathcal{I}'} = \{0\}$ then $\mathcal{U} \cup \{\delta_j : j \in \mathcal{I}'\}$ forms a basis for \mathbb{C}^N . For this, let $f \in \mathbb{Y} \cap \mathbb{C}^{\mathcal{I}'}$. Then

$$f = \sum_{i \in \mathcal{I}} f(i)u_i \quad (2)$$

because \mathcal{U} is an interpolating basis for \mathbb{Y} , and also

$$f = \sum_{j \in \mathcal{I}'} f(j)\delta_j.$$

Thus

$$\sum_{i \in \mathcal{I}} f(i)u_i = \sum_{j \in \mathcal{I}'} f(j)\delta_j.$$

Let $k \in \mathcal{I}$ and evaluate both sides at k :

$$\begin{aligned} \sum_{i \in \mathcal{I}} f(i)u_i(k) &= \sum_{j \in \mathcal{I}'} f(j)\delta_j(k), \\ f(k) &= 0. \end{aligned}$$

By (2), $f = 0$ and we are done. ■

The algebraic proof of Theorem 1 is in terms of matrices. Associate with an index set $\mathcal{I} = \{i_1, i_2, \dots, i_d\}$ the $N \times d$ matrix $E_{\mathcal{I}}$ whose d columns are the basis vectors $\delta_{i_1}, \delta_{i_2}, \dots, \delta_{i_d}$. If R is an $N \times M$ matrix then $E_{\mathcal{I}}^T R$ is $d \times M$ submatrix of R obtained by choosing the rows indexed by \mathcal{I} . In particular, operating by $E_{\mathcal{I}}^T$ on an N -vector f produces the d -vector with components $f(i_1), f(i_2), \dots, f(i_d)$. If R is an $M \times N$ matrix then $RE_{\mathcal{I}}$ is the $M \times d$ submatrix of R obtained by choosing the columns indexed by \mathcal{I} .

We note three general facts. First, $E_{\mathcal{I}}^T E_{\mathcal{I}} = I_d$, where I_d is the $d \times d$ identity matrix. Second, if S is a $d \times d$ matrix then $E_{\mathcal{I}}^T (RS) = (E_{\mathcal{I}}^T R)S$. Finally, if $\mathcal{U} = \{u_{i_1}, u_{i_2}, \dots, u_{i_d}\}$ is a basis for \mathbb{Y} and U is the $N \times d$ matrix whose columns are the u_i then the condition (1) that \mathcal{U} be an interpolating basis can be written in matrix form as

$$f = UE_{\mathcal{I}}^T f \quad (3)$$

for all $f \in \mathbb{Y}$. Here $UE_{\mathcal{I}}^T$ is an $N \times N$ matrix and we see that \mathcal{U} is an interpolating basis for \mathbb{Y} with sampling set \mathcal{I} if and only if $\mathbb{Y} = \ker(I_N - UE_{\mathcal{I}}^T)$.

Now we have

Algebraic Proof of Theorem 1: Take any basis $\mathcal{V} = \{v_1, v_2, \dots, v_d\}$ of \mathbb{Y} and let R be the $N \times d$ matrix whose columns are the basis vectors v_k ; thus $R_{jk} = v_k(j)$. Since R has rank d it has a $d \times d$ invertible submatrix, and possibly many such submatrices. Let \mathcal{I} be the index set corresponding to the d rows chosen from R to form the invertible submatrix $E_{\mathcal{I}}^T R$. The columns of the $N \times d$ matrix $R(E_{\mathcal{I}}^T R)^{-1}$ are again

a basis of \mathbb{Y} . We write them as $u_{i_1}, u_{i_2}, \dots, u_{i_d}$, indexed by \mathcal{I} . Since

$$E_{\mathcal{I}}^T (R(E_{\mathcal{I}}^T R)^{-1}) = (E_{\mathcal{I}}^T R)(E_{\mathcal{I}}^T R)^{-1} = I_d,$$

the u_{i_j} are as in Proposition 1, and hence comprise an interpolating basis of \mathbb{Y} . ■

This proof shows how to produce an interpolating basis provided one can find a $d \times d$ invertible submatrix $E_{\mathcal{I}}^T R$, indexed by \mathcal{I} . The more such submatrices the more interpolating bases for \mathbb{Y} . On the opposite side, in general not every index set \mathcal{I} is sampling set for an interpolating basis since, in general, not every choice of a $d \times d$ submatrix is invertible.

A slightly different way of arranging the algebraic proof also gives an interpolation formula, making (3) more explicit. As above, let $\mathcal{V} = \{v_1, v_2, \dots, v_d\}$ be a basis of \mathbb{Y} and let R be the corresponding $N \times d$ matrix. If $f \in \mathbb{Y}$ then

$$f = \sum_{k=1}^N f(k)\delta_k \quad \text{and also} \quad f = \sum_{k=1}^d \alpha_k v_k,$$

for some constants α_k . We want to solve for the α_k in terms of d of the values $f(k)$. Write the second equation for f as

$$f = R\alpha, \quad \alpha = (\alpha_1, \alpha_2, \dots, \alpha_d)^T.$$

Now R has an invertible $d \times d$ submatrix, say $E_{\mathcal{I}}^T R$ for an index set \mathcal{I} , and so

$$E_{\mathcal{I}}^T f = E_{\mathcal{I}}^T (R\alpha) = (E_{\mathcal{I}}^T R)\alpha.$$

We can then solve for α via

$$\alpha = (E_{\mathcal{I}}^T R)^{-1} (E_{\mathcal{I}}^T f),$$

resulting in

$$f = R(E_{\mathcal{I}}^T R)^{-1} (E_{\mathcal{I}}^T f). \quad (4)$$

This equation writes f in terms of the components $f(i)$, $i \in \mathcal{I}$.

Carrying the algebraic line of reasoning a little further, we also see how two interpolating bases for \mathbb{Y} are related to each other.

Theorem 2: Fix an interpolating basis of \mathbb{Y} , indexed by \mathcal{J} , and let R be the corresponding $N \times d$ matrix. If S is the matrix of another interpolating basis of \mathbb{Y} , indexed by \mathcal{I} , then $E_{\mathcal{I}}^T R$ is invertible and

$$S = R(E_{\mathcal{I}}^T R)^{-1}.$$

Proof: Let $\{v_i : i \in \mathcal{I}\}$ be the interpolating basis of \mathbb{Y} that are the columns of S and let $\{u_j : j \in \mathcal{J}\}$ be the columns of R . Since the u_j are an interpolating basis we can write, for each $i \in \mathcal{I}$,

$$v_i = \sum_{j \in \mathcal{J}} v_i(j)u_j.$$

In matrix form this is

$$S = R(E_{\mathcal{I}}^T S).$$

Now multiply on the left by $E_{\mathcal{I}}^T$, resulting in

$$E_{\mathcal{I}}^T S = E_{\mathcal{I}}^T (R(E_{\mathcal{I}}^T S)) = (E_{\mathcal{I}}^T R)(E_{\mathcal{I}}^T S).$$

But $E_{\mathcal{I}}^T R$ is the $d \times d$ identity matrix, so this shows that $E_{\mathcal{I}}^T R$ is invertible, that $(E_{\mathcal{I}}^T R)^{-1} = E_{\mathcal{I}}^T S$, and then that

$$S = R(E_{\mathcal{I}}^T R)^{-1}.$$

Finally, we look a little more closely at the interpolating basis provided by $R(E_{\mathcal{I}}^T R)^{-1}$ in relation to the geometric construction. From the $d \times d$ matrix $(E_{\mathcal{I}}^T R)^{-1}$ form a $d \times N$ matrix by adding $N - d$ columns of zeros in the slots \mathcal{I}' . Call this matrix T . Then RT is an $N \times N$ matrix and one sees that

$$RT\delta_i = \begin{cases} u_i, & i \in \mathcal{I} \\ 0, & i \in \mathcal{I}' \end{cases}$$

Thus RT is the projection of \mathbb{C}^N onto \mathbb{Y} along $\mathbb{C}^{\mathcal{I}'}$ and we are back to the idea of the geometric argument. Observe that whereas the geometric argument started with a complement $\mathbb{C}^{\mathcal{I}'}$ to \mathbb{Y} and produced the interpolating basis via projection, here we started with an interpolating basis for \mathbb{Y} and produced the projection and the complement.

III. DISCRETE BANDLIMITED SPACES

Bandlimited signals are defined by the vanishing of the discrete Fourier transform outside a set of specified indices. They form a particularly interesting class of subspaces.

For notation, let

$$\zeta_n = e^{-2\pi i/n},$$

simplified to just ζ when $n = N$, and let $\omega: \mathbb{Z}_N \rightarrow \mathbb{C}$ be the discrete complex exponential,

$$\omega(m) = \zeta^m.$$

The discrete Fourier transform is then

$$\mathcal{F}f = \sum_{n=0}^{N-1} f(n)\omega^n.$$

As usual, we also regard \mathcal{F} as an $N \times N$ matrix whose mn -entry is $\mathcal{F}_{mn} = \omega^n(m) = \zeta^{mn}$. We recall that $\mathcal{F}^{-1} = (1/N)\mathcal{F}^*$ (the adjoint of \mathcal{F}).

Definition 2: Let $\mathcal{J} \subseteq [0 : N - 1]$. The $|\mathcal{J}|$ -dimensional space of bandlimited signals with frequency support \mathcal{J} is

$$\mathbb{B}^{\mathcal{J}} = \mathcal{F}^{-1}(\mathbb{C}^{\mathcal{J}}).$$

In words, $f \in \mathbb{B}^{\mathcal{J}}$ if $\mathcal{F}f$ has zeros in the slots $\mathcal{J}' = [1 : N] \setminus \mathcal{J}$. There might be more zeros of $\mathcal{F}f$ for a given f but there are at least these zeros. We do not assume that the indices in \mathcal{J} are contiguous, so $\mathcal{F}f$ is not necessarily supported on a “band” of frequencies, but we maintain the use of the term “bandlimited” in all cases.

Since $\mathcal{F}^*\delta_n(m) = \zeta^{-mn}$, we get a basis for $\mathbb{B}^{\mathcal{J}}$ by pulling out of \mathcal{F}^* the columns indexed by \mathcal{J} . Thus we get an interpolating basis with sampling set \mathcal{I} if and only if $E_{\mathcal{I}}^T \mathcal{F}^* E_{\mathcal{J}}$ is invertible, or equivalently if and only if $E_{\mathcal{I}}^T \mathcal{F} E_{\mathcal{J}}$ is invertible. We prefer to use the latter, with \mathcal{F} instead of \mathcal{F}^* .

For the remainder of this paper, interpolating systems for bandlimited spaces will be our main concern. Spaces of bandlimited functions having orthogonal interpolating bases are the subject of [1], but we do have one general observation here: such spaces cannot be too big.

Proposition 4: If $\mathbb{B}^{\mathcal{J}}$ has an orthogonal interpolating basis then $|\mathcal{J}| \leq N/2$.

Proof: Suppose $\mathbb{B}^{\mathcal{J}}$ has an orthogonal interpolating basis indexed by \mathcal{I} . Then $|\mathcal{I}| = |\mathcal{J}|$. Let $\mathcal{I}' = [0 : N - 1] \setminus \mathcal{I}$. By Proposition 2 we can write

$$\mathbb{B}^{\mathcal{J}} = \text{span}\{\delta_i + v_i : i \in \mathcal{I}\},$$

where the v_i are orthogonal vectors in $\mathbb{C}^{\mathcal{I}'}$, or some possibly 0. But none of the v_i can be zero, for $\mathcal{F}\delta_i = \omega^{-k}$ which never vanishes. There are $|\mathcal{I}|$ of the v 's, and if $|\mathcal{J}| = |\mathcal{I}| > N/2$ then $|\mathcal{I}'| < N/2$ and we would have more than $N/2$ orthogonal vectors in a space of dimension less than $N/2$. ■

A. Necklaces and Bracelets

Sampling sets for bandlimited spaces have more algebraic structure than it might appear. Namely, the property of being a sampling set for a particular $\mathbb{B}^{\mathcal{J}}$ is preserved under the action of the dihedral group. To explain, on \mathbb{Z}_N we denote the operations of translation (by 1) and reflection by τ and ρ , respectively:

$$\begin{aligned} \tau: \mathbb{Z}_N &\rightarrow \mathbb{Z}_N, & \tau(n) &= n - 1 \pmod{N}, \\ \rho: \mathbb{Z}_N &\rightarrow \mathbb{Z}_N, & \rho(n) &= -n \pmod{N}. \end{aligned}$$

Then

$$\tau^N = \text{id}, \quad \rho^2 = \text{id} \quad \text{and} \quad \rho\tau\rho = \tau^{-1} \quad \text{or} \quad (\rho\tau)^2 = \text{id},$$

so τ and ρ generate the dihedral group Dih_N . Clearly Dih_N can act on an index set \mathcal{I} via

$$\tau\mathcal{I} = \{\tau(i) : i \in \mathcal{I}\}, \quad \rho\mathcal{I} = \{\rho(i) : i \in \mathcal{I}\}.$$

We define the *bracelet* of \mathcal{I} to be the orbit of \mathcal{I} under the action of Dih_N . The *necklace* of \mathcal{I} is the orbit of \mathcal{I} under the action of the cyclic subgroup $\langle \tau \rangle$ of Dih_N . Think of $\mathcal{I} \subset [0 : N - 1]$ as specifying a pattern of N beads on a loop, with black beads in the locations in \mathcal{I} separated by white beads in the locations in the complement \mathcal{I}' , as in Figure 1. A necklace is worn around the neck, and if the cyclic group acts then the spacing of the black and white beads is the same however the necklace is rotated. But a bracelet can be worn on either wrist, introducing a reflection, and the symmetry group is Dih_N . See Appendix B for a formula that counts distinct bracelets, and for references.

With these definitions we now have

Proposition 5: If \mathcal{I} is a sampling set for $\mathbb{B}^{\mathcal{J}}$ then any index set in the bracelet of \mathcal{I} is a sampling set for $\mathbb{B}^{\mathcal{J}}$.

Proof: Let $\mathcal{I} = \{m_1, m_2, \dots, m_d\}$, $\mathcal{J} = \{n_1, n_2, \dots, n_d\}$ and let $\mathcal{K} = \tau\mathcal{I}$. Then the new submatrix $E_{\mathcal{K}}^T \mathcal{F} E_{\mathcal{J}}$ is given by

$$\begin{aligned} E_{\mathcal{K}}^T \mathcal{F} E_{\mathcal{J}} &= \begin{bmatrix} \zeta^{(m_1-1)n_1} & \zeta^{(m_1-1)n_2} & \dots & \zeta^{(m_1-1)n_d} \\ \zeta^{(m_2-1)n_1} & \zeta^{(m_2-1)n_2} & \dots & \zeta^{(m_2-1)n_d} \\ \vdots & \vdots & \ddots & \vdots \\ \zeta^{(m_d-1)n_1} & \zeta^{(m_d-1)n_2} & \dots & \zeta^{(m_d-1)n_d} \end{bmatrix} \\ &= \begin{bmatrix} \zeta^{m_1 n_1} & \zeta^{m_1 n_2} & \dots & \zeta^{m_1 n_d} \\ \zeta^{m_2 n_1} & \zeta^{m_2 n_2} & \dots & \zeta^{m_2 n_d} \\ \vdots & \vdots & \ddots & \vdots \\ \zeta^{m_d n_1} & \zeta^{m_d n_2} & \dots & \zeta^{m_d n_d} \end{bmatrix} \times \end{aligned}$$




Fig. 1: Two different bracelets with $N = 12$ and $|\mathcal{I}| = 4$. On top the index set is $\mathcal{I} = \{0, 2, 5, 7\}$, on the bottom the index set is $\mathcal{I} = \{0, 3, 5, 6\}$

$$\begin{bmatrix} \zeta^{-n_1} & 0 & 0 & \cdots & 0 \\ 0 & \zeta^{-n_2} & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \zeta^{-n_d} \end{bmatrix} \\
 = (E_{\mathcal{I}}^T \mathcal{F} E_{\mathcal{J}}) \times \text{an invertible diagonal matrix.}$$

Hence $E_{\mathcal{K}}^T \mathcal{F} E_{\mathcal{J}}$ is invertible whenever $E_{\mathcal{I}}^T \mathcal{F} E_{\mathcal{J}}$ is, and the same is true for any translation of \mathcal{I} .

Next suppose \mathcal{K} is obtained by reversing \mathcal{I} , namely $\mathcal{K} = \{N - m_1, N - m_2, \dots, N - m_d\}$. Then $E_{\mathcal{K}}^T \mathcal{F} E_{\mathcal{J}}$ is just the conjugate of $E_{\mathcal{I}}^T \mathcal{F} E_{\mathcal{J}}$, so again, $E_{\mathcal{K}}^T \mathcal{F} E_{\mathcal{J}}$ is invertible whenever $E_{\mathcal{I}}^T \mathcal{F} E_{\mathcal{J}}$ is. ■

IV. UNIVERSAL SAMPLING SETS

There is a kind of interchange duality for bandlimited spaces between sampling sets and frequency support sets. On the one hand, the sampling problem is to start with $\mathbb{B}^{\mathcal{J}}$ and ask which index sets \mathcal{I} are sampling sets. On the other hand, one could also start with an index set \mathcal{I} and ask which $\mathbb{B}^{\mathcal{J}}$ result from this sampling pattern. These two questions are equivalent.

Proposition 6: $\mathbb{B}^{\mathcal{J}}$ has \mathcal{I} as a sampling set if and only if $\mathbb{B}^{\mathcal{I}}$ has \mathcal{J} as a sampling set.

Proof: The subspace $\mathbb{B}^{\mathcal{J}}$ has \mathcal{I} as a sampling set if and only if $E_{\mathcal{I}}^T \mathcal{F} E_{\mathcal{J}}$ is invertible, and this is true if and only if its transpose $E_{\mathcal{J}}^T \mathcal{F} E_{\mathcal{I}}$ is invertible. ■

Though the sampling problem may seem the more natural one, we will concentrate on the second, equivalent question and ask which frequency patterns, that is which $\mathbb{B}^{\mathcal{J}}$, can arise

from a given sampling set \mathcal{I} . It may be that the space $\mathbb{B}^{\mathcal{J}}$ is not known exactly, or that we may have some erroneous estimate $\tilde{\mathcal{J}}$ of \mathcal{J} . The question is whether we can pick sampling locations \mathcal{I} that are robust for these estimation errors. We will find some interesting phenomena, and the results can easily be translated to apply to the sampling problem. The extreme case is captured by the following definition.

Definition 3: An index set $\mathcal{I} \subset [0 : N - 1]$ is a *universal sampling set* if \mathcal{I} is a sampling set for each $\mathbb{B}^{\mathcal{J}}$ with $|\mathcal{J}| = |\mathcal{I}|$. See also [4] and [5].

If \mathcal{I} is a universal sampling set, then while an interpolating basis of a space $\mathbb{B}^{\mathcal{J}}$ still depends on \mathcal{J} , where the samples are taken does *not* depend on \mathcal{J} . In Section V we will show that there are universal sampling sets of any given size; in fact, we will count them.

Very concretely, to ask if \mathcal{I} is a universal sampling set is to ask if there are rows of \mathcal{F} indexed by \mathcal{I} , $|\mathcal{I}| = d$, such that *any* $d \times d$ submatrix of \mathcal{F} formed with these rows is invertible. Phrased this way, standard properties of Vandermonde determinants applied to \mathcal{F} allow us to conclude:

Proposition 7: (i) If \mathcal{I} is a set of d consecutive indices, reduced mod N ,

$$\mathcal{I} = \{i_0, i_0 + 1, \dots, i_0 + (d - 1)\} \pmod{N},$$

then \mathcal{I} is a universal sampling set.

(ii) If \mathcal{I} is a set of d indices in arithmetic progression, reduced mod N ,

$$\mathcal{I} = \{i_0, i_0 + s, i_0 + 2s, \dots, i_0 + (d - 1)s\} \pmod{N},$$

where s is coprime to N , then \mathcal{I} is a universal sampling set. ■

Much deeper is the following theorem of Chebotarev.

Theorem 3: (Chebotarev) If N is prime, then every square submatrix of \mathcal{F} is invertible.

And so, if N is prime then any index set \mathcal{I} is a universal sampling set. Chebotarev's theorem dates to 1948 (the original paper is in Russian) and there are now several published (and unpublished) proofs, see, e.g., [8], [9], but this is by no means a trivial result.

We will generalize Chebotarev's theorem when N is a prime power, and we will offer several characteristic properties of universal sampling sets. We are indebted to the works of Tao [6] and Delvaux and Van Barel [10].

The key is a quantitative, almost statistical comparison of \mathcal{I} to the simplest universal sampling set,

$$\mathcal{I}^* = [0 : d - 1],$$

when the elements of both \mathcal{I} and \mathcal{I}^* are reduced modulo prime powers. We need several additional definitions to state our main results.

A. Multisets and the Size of Congruence Classes

We have found it conceptually helpful to use *multisets* in the description of one of the central ideas, and we briefly review this concept. Informally, a multiset is a finite, unordered list \tilde{A} whose elements are drawn from a finite set A , and

where, to distinguish a multiset from simply a set, elements of the list may be repeated. More formally, a multiset is a pair $(A, \tilde{\chi}_A)$ where $\tilde{\chi}_A$ is the multiplicity function (generalizing the characteristic function):

$$\tilde{\chi}_A: A \longrightarrow \mathbb{N},$$

$$\tilde{\chi}_A(a) = \text{the number of times } a \in A \text{ is listed in } \tilde{A}.$$

Two multisets \tilde{A} and \tilde{B} are equal if $\tilde{\chi}_A = \tilde{\chi}_B$, so the individual elements are the same and so are their multiplicities. The cardinality of \tilde{A} is

$$|\tilde{A}| = \sum_{a \in A} \tilde{\chi}_A(a).$$

It is common practice to use the standard set notation in writing a multiset. Thus, for example, drawing from $\{a, b, c, d\}$ we write a multiset as $\{a, a, c, d, d\}$. The tilde notation \tilde{A} for a multiset drawn from A is helpful in discussing general principles but, like all general notations, it has its limitations in particular cases. It is a notation often used for covering spaces, as we comment on below.

Associated with a multiset \tilde{A} is another multiset

$$\mathcal{M}(\tilde{A}) = \{\tilde{\chi}_A(a) : a \in A\},$$

which we call the *multiplicity multiset* of \tilde{A} . Thus $\mathcal{M}(\tilde{A})$ records as a multiset the counts of the elements of \tilde{A} and also includes a zero for each element of A that does not appear in \tilde{A} . One can think of $\mathcal{M}(\tilde{A})$ as providing some statistics of \tilde{A} , a kind of histogram of \tilde{A} with bins from A , except that the bins are not ordered.

Next, let p be a prime, $k \geq 0$ an integer, and for $x \in \mathbb{N}$ let $[x]_k$ be the residue of x reduced mod p^k . For an index set \mathcal{I} let

$$\mathcal{I}/p^k = \{[i]_k : i \in \mathcal{I}\}$$

be the set of residues mod p^k of the elements of \mathcal{I} , and let $(\mathcal{I}/p^k)^\sim$ be the corresponding multiset, meaning that each residue is listed according to its multiplicity, i.e., the size of its congruence class. We regard the elements of $(\mathcal{I}/p^k)^\sim$ to be drawn from $[0 : p^k - 1]$, all possible residues, and we write $\tilde{\chi}_k: [0 : p^k - 1] \longrightarrow \mathbb{N}$ for the multiplicity function for the multiplicity multiset $\mathcal{M}((\mathcal{I}/p^k)^\sim)$. To be explicit, for $a \in [0 : p^k - 1]$

$$\tilde{\chi}_k(a) = \text{the number of elements of } \mathcal{I} \text{ that leave a remainder of } a \text{ on dividing by } p^k. \quad (5)$$

In particular, $\tilde{\chi}_k(a) = 0$ means that no element of \mathcal{I} leaves a remainder of a on dividing by p^k . In this case we speak of an empty congruence class in \mathcal{I}/p^k . For $a \in [0 : p^k - 1]$ it will also be helpful to use the notation

$$\mathcal{I}_{ka} = \{i \in \mathcal{I} : i \equiv a \pmod{p^k}\}$$

for the elements of the congruence class of $a \pmod{p^k}$. Then $\tilde{\chi}_k(a) = |\mathcal{I}_{ka}|$.

When we need to emphasize the index set, especially in Section V, we will write $\tilde{\chi}_k(a; \mathcal{I})$. We note the obvious properties:

- If \mathcal{I} and \mathcal{J} are disjoint then $\tilde{\chi}_k(a; \mathcal{I} \cup \mathcal{J}) = \tilde{\chi}_k(a; \mathcal{I}) + \tilde{\chi}_k(a; \mathcal{J})$.

$$\bullet \quad \mathcal{I} \subseteq \mathcal{J} \implies \tilde{\chi}_k(a; \mathcal{I}) \leq \tilde{\chi}_k(a; \mathcal{J}).$$

Observe for $k = 0$ that $(\mathcal{I}/1)^\sim$ just consists of $|\mathcal{I}|$ zeros and $\tilde{\chi}_0(0) = |\mathcal{I}|$. More generally,

$$|\mathcal{I}| = \sum_{a=0}^{p^k-1} \tilde{\chi}_k(a). \quad (6)$$

We also note that the multiplicity multiset $\mathcal{M}((\mathcal{I}/p^k)^\sim)$ depends only on the bracelet of \mathcal{I} . While the multisets $(\mathcal{I}/p^k)^\sim$ will generally change if \mathcal{I} is shifted or reversed, the counts of the residues on dividing by p^k will be the same:

$$\mathcal{M}((\tau\mathcal{I}/p^k)^\sim) = \mathcal{M}((\mathcal{I}/p^k)^\sim) \quad \text{and} \quad \mathcal{M}((\rho\mathcal{I}/p^k)^\sim) = \mathcal{M}((\mathcal{I}/p^k)^\sim). \quad (7)$$

Remark 1: Introducing the multiset $(\mathcal{I}/p^k)^\sim$ is reminiscent of introducing covering spaces (for Riemann surfaces) to resolve the problem of multivalued functions. Here we have the remainder map $r: \mathcal{I} \longrightarrow \mathcal{I}/p^k$, $r(i) = [i]$, which is generally not injective and so has a multivalued inverse. Think of the residues (with multiplicity) in $(\mathcal{I}/p^k)^\sim$ as tagged by the number they come from, say as a pair $([i], i)$, which serves to distinguish them much as we think of tagging points on different sheets of a covering space of a Riemann surface. Then we have the commutative diagram

$$\begin{array}{ccc} & (\mathcal{I}/p^k)^\sim & \\ \tilde{r} \nearrow & \downarrow \text{pr} & \\ \mathcal{I} & \xrightarrow{r} & \mathcal{I}/p^k \end{array}$$

where pr is the projection map, $([i], i) \mapsto [i]$ and the lift $\tilde{r}(i) = ([i], i)$, of r is bijective. The value of the multiplicity function $\tilde{\chi}_k(i)$ is then the number of elements in the preimage $\text{pr}^{-1}([i])$, analogous to the number of sheets over $[i]$. It will generally vary with $[i]$.

Returning to our primary considerations, we write $\tilde{\chi}_k^*$ to distinguish the special case when $\mathcal{I} = \mathcal{I}^*$. We will need the following property of $\tilde{\chi}_k^*$:

$$|\tilde{\chi}_k^*(a) - \tilde{\chi}_k^*(b)| \leq 1, \quad (8)$$

for all $a, b \in [0 : p^k - 1]$ and all k . In words, when reducing the elements of $\mathcal{I}^* = [0 : d - 1]$ modulo p^k for any k , the conjugacy classes are all of about the same size. Or, pursuing the analogy above, the preimages $\text{pr}^{-1}([i])$ of the individual residues all have approximately the same number of elements and one might say that \mathcal{I}^*/p^k is *uniformly covered* for each k .

The inequality in (8) is easy to see. For some background calculations we have found it helpful to have a formula for $\tilde{\chi}_k^*$ (from which (8) also follows). If $\ell \in \mathcal{I}^*$ with $[\ell]_k = a \in [0 : p^k - 1]$ then $\ell = a + \alpha p^k$ for an integer $\alpha \geq 0$, and since $\ell \leq d - 1$ we must have $0 \leq \alpha \leq (d - 1 - a)/p^k$. The number of integers α for which this inequality holds is the number of ℓ whose residue is a . Thus

$$\tilde{\chi}_k^*(a) = \left\lfloor \frac{d - 1 - a}{p^k} + 1 \right\rfloor. \quad (9)$$

B. A Characterization of Universal Sampling Sets

Our main result is:

Theorem 4: Let \mathcal{I} be an index set in $[0 : p^M - 1]$. The following are equivalent:

- (i) $\tilde{\chi}_k = \tilde{\chi}_k^*$ for all $0 \leq k \leq M$.
- (ii) $|\tilde{\chi}_k(a) - \tilde{\chi}_k(b)| \leq 1$ for all $a, b \in [0 : p^k - 1]$ and $0 \leq k \leq M$.
- (iii) \mathcal{I} is a universal sampling set.

According to Proposition 5 and the relations (7), any index set in the bracelet of \mathcal{I} is also a universal sampling set. Likewise, any index set in the bracelet of \mathcal{I}^* can serve as a model universal sampling set. Only condition (i) directly compares \mathcal{I} to \mathcal{I}^* , and in terms of multisets it could be stated equivalently as

$$\mathcal{M}((\mathcal{I}/p^k)^\sim) = \mathcal{M}((\mathcal{I}^*/p^k)^\sim).$$

Condition (i) for $k = 0$ guarantees that \mathcal{I} and \mathcal{I}^* have the same size, from (6). Computing $\mathcal{M}((\mathcal{I}/p^k)^\sim)$ and $\mathcal{M}((\mathcal{I}^*/p^k)^\sim)$ for $k \geq M$ is redundant; since all elements in \mathcal{I} and \mathcal{I}^* are in $[0 : p^M - 1]$, $\mathcal{M}((\mathcal{I}/p^k)^\sim)$ for $k \geq M$ is just indicative of the cardinality of \mathcal{I} and \mathcal{I}^* . Namely, for $k \geq M$, each of $\mathcal{M}((\mathcal{I}/p^k)^\sim)$ and $\mathcal{M}((\mathcal{I}^*/p^k)^\sim)$ contains $|\mathcal{I}|$ ones and $p^k - |\mathcal{I}|$ zeros. Condition (ii), a property only of \mathcal{I} , indirectly compares \mathcal{I} to \mathcal{I}^* via (8). It says that \mathcal{I}/p^k , like \mathcal{I}^*/p^k , is uniformly covered for each k .

Before we embark on the proof of the theorem, here is an example. Let $N = 2^3$, and $\mathcal{I} = \{0, 1, 3, 4, 6\}$. The following are the multisets for $k = 1, 2, 3$:

$$\begin{aligned} (\mathcal{I}/2)^\sim &= \{0, 1, 1, 0, 0\}, & \mathcal{M}((\mathcal{I}/2)^\sim) &= \{3, 2\}; \\ (\mathcal{I}/2^2)^\sim &= \{0, 1, 3, 0, 2\}, & \mathcal{M}((\mathcal{I}/2^2)^\sim) &= \{2, 1, 1, 1\}; \\ (\mathcal{I}/2^3)^\sim &= \{0, 1, 3, 4, 6\}, \\ \mathcal{M}((\mathcal{I}/2^3)^\sim) &= \{1, 1, 0, 1, 1, 0, 1, 0\}. \end{aligned}$$

The computations for $\mathcal{I}^* = \{0, 1, 2, 3, 4\}$ yield

$$\begin{aligned} (\mathcal{I}^*/2)^\sim &= \{0, 1, 0, 1, 0\}, & \mathcal{M}((\mathcal{I}^*/2)^\sim) &= \{3, 2\}; \\ (\mathcal{I}^*/2^2)^\sim &= \{0, 1, 2, 3, 0\}, & \mathcal{M}((\mathcal{I}^*/2^2)^\sim) &= \{2, 1, 1, 1\}; \\ (\mathcal{I}^*/2^3)^\sim &= \{0, 1, 2, 3, 4\}, \\ \mathcal{M}((\mathcal{I}^*/2^3)^\sim) &= \{1, 1, 1, 1, 1, 0, 0, 0\}. \end{aligned}$$

We see that $\mathcal{M}((\mathcal{I}/2^k)^\sim) = \mathcal{M}((\mathcal{I}^*/2^k)^\sim)$ for $k = 1, 2, 3$, and hence \mathcal{I} is a universal sampling set. So in case the reader has ever wondered, for the 8×8 Fourier matrix any 5×5 submatrix built from the rows indexed by \mathcal{I} , or from the rows of an index set in the bracelet of \mathcal{I} , is invertible.

Proof of Theorem 4, (i) \iff (ii): Note: This equivalence does not require that N be a prime power. The implication (i) \implies (ii) is immediate from (8). Assume (ii) holds and let

$$\chi = \min_a \tilde{\chi}_k(a).$$

From (ii) it follows that any $\tilde{\chi}_k(a)$ is either χ or $\chi + 1$. Suppose r of the p^k numbers $\tilde{\chi}_k(a)$ are equal to $\chi + 1$ and the rest are equal to χ . The cardinality equation, (6),

$$\sum_{a=0}^{p^k-1} \tilde{\chi}_k(a) = |\mathcal{I}| = d, \quad (10)$$

then gives

$$p^k \chi + r = d, \quad \text{with} \quad 0 \leq r < p^k.$$

This means that χ is the quotient on dividing d by p^k and r is the remainder. In other words, (ii) and (10) together uniquely determine the multiset $\mathcal{M}((\mathcal{I}/p^k)^\sim) = \{\tilde{\chi}_k(a) : a \in [0, p^k - 1]\}$. Since \mathcal{I} and \mathcal{I}^* both satisfy (ii) and (10), we must have $\mathcal{M}((\mathcal{I}/p^k)^\sim) = \mathcal{M}((\mathcal{I}^*/p^k)^\sim)$, or $\tilde{\chi}_k = \tilde{\chi}_k^*$. \blacksquare

We need two lemmas to prove that condition (i) implies that \mathcal{I} is a universal sampling set. The first is a very old theorem on Vandermonde determinants, [11], as updated in [10]:

Lemma 1 (Delvaux and Van Barel): Let

$$V = \begin{bmatrix} x_1^{m_1} & x_2^{m_1} & x_3^{m_1} & \cdots & x_d^{m_1} \\ x_1^{m_2} & x_2^{m_2} & x_3^{m_2} & \cdots & x_d^{m_2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ x_1^{m_d} & x_2^{m_d} & x_3^{m_d} & \cdots & x_d^{m_d} \end{bmatrix} \quad (11)$$

be a $d \times d$ generalized Vandermonde matrix. Then the determinant of V is given by

$$\det V = \left(\prod_{i < j} (x_j - x_i) \right) S(x_1, x_2, \dots, x_d), \quad (12)$$

where $S(x_1, x_2, \dots, x_d)$ is a symmetric polynomial in x_1, x_2, \dots, x_d with integer coefficients such that

$$S(1, 1, \dots, 1) = \frac{\prod_{0 \leq i < j \leq d-1} (m_j - m_i)}{\prod_{0 \leq i < j \leq d-1} (j - i)}.$$

The polynomial S is called a *Schur polynomial*, see, for example, [12]. Based on this lemma we deduce a second result that is itself already a sufficient condition for an index set to be a universal sampling set.

Lemma 2: Let $\mathcal{I} = \{m_0, m_1, m_2, \dots, m_{d-1}\}$. If

$$\mu = \frac{\prod_{0 \leq i < j \leq d-1} (m_j - m_i)}{\prod_{0 \leq i < j \leq d-1} (j - i)} \quad (13)$$

is coprime to p , then \mathcal{I} is a universal sampling set.

Note that without Lemma 1, it would not even be clear that μ is an integer. An intuitive idea for why this should be so is given below. The proof of Lemma 2 is along the lines of the proof of Chebotarev's theorem in [13], and also in [6].

Proof of Lemma 2: We make use of Lemma 1 in the case when $V = E_{\mathcal{I}}^T F E_{\mathcal{J}}$. Each x_ℓ in (11) is then a power of $\zeta = e^{-2\pi i/N}$, $x_\ell = \zeta^{j_\ell}$, where $\mathcal{J} = \{j_1, j_2, \dots, j_d\}$.

Suppose $\det V = 0$. From (12), this means that $S(x_1, x_2, \dots, x_d) = 0$. Substituting $x_\ell = \zeta^{j_\ell}$ in $S(x_1, x_2, \dots, x_d) = 0$, we obtain an equation of the form $s(\zeta) = 0$, where $s(x)$ is a polynomial in one variable with integer coefficients. This means that ζ is a root of $s(x)$ and since $s(x)$ has only integer coefficients, $s(x)$ must contain the minimal polynomial of ζ over \mathbb{Z} as a factor.

For $N = p^M$, the minimal polynomial of ζ over \mathbb{Z} is $\phi_N(x) = 1 + x^{p^{M-1}} + x^{2p^{M-1}} + x^{3p^{M-1}} + \dots + x^{(p-1)p^{M-1}}$ (the N 'th cyclotomic polynomial). So we have $\phi_N(x) \mid s(x)$, where

$$\phi_N(x) = 1 + x^{p^{M-1}} + x^{2p^{M-1}} + x^{3p^{M-1}} + \dots + x^{(p-1)p^{M-1}}.$$

Now ϕ_N and s are both polynomials with integer coefficients, hence $\phi_N(1) \mid s(1)$. However, $\phi_N(1) = p$, and $s(1) = S(1, 1, \dots, 1) = \mu$. Thus

$$p \mid \mu \text{ if } \det V = 0.$$

This proves the lemma. \blacksquare

Chebotarev's theorem follows from this result. If N is a prime p then μ is coprime to p because every factor in the numerator and denominator of μ is an integer strictly between $-p$ and p .

We can now complete the proof of one direction of the implications in Theorem 4.

Proof of Theorem 4: (i) \implies (iii): Let $\mathcal{I} = \{m_1, m_2, m_3, \dots, m_d\}$ and consider the product of differences

$$A = \prod_{1 \leq i < j \leq d} (m_j - m_i).$$

There are $\tilde{\chi}_k(\ell)$ elements of \mathcal{I} that leave a remainder of ℓ when divided by p^k . Moreover, $m_i \equiv m_j \pmod{p^k}$ if and only if $p^k \mid (m_j - m_i)$. The number of differences that have a factor of p^k (or higher: p^r for $r > k$) is

$$\sum_{l=0}^{p^k-1} \binom{\tilde{\chi}_k(l)}{2},$$

and hence the number of differences that have a factor of exactly p^k is given by

$$\sum_{l=0}^{p^k-1} \binom{\tilde{\chi}_k(l)}{2} - \sum_{l=0}^{p^{k+1}-1} \binom{\tilde{\chi}_{k+1}(l)}{2}.$$

The largest power of p that divides A is then p raised to

$$\sum_k k \left(\sum_{l=0}^{p^k-1} \binom{\tilde{\chi}_k(l)}{2} - \sum_{l=0}^{p^{k+1}-1} \binom{\tilde{\chi}_{k+1}(l)}{2} \right). \quad (14)$$

The expression (14) depends only on the values of $\tilde{\chi}_k$, but the hypothesis is that $\tilde{\chi}_k = \tilde{\chi}_k^*$ for $0 \leq k \leq N$, and therefore the products $A = \prod (m_j - m_i)$ and $B = \prod (j - i)$ have the same powers of p as factors. Hence $\mu = A/B$ is coprime to p and from Lemma 2 we conclude that \mathcal{I} is a universal sampling set. \blacksquare

Remark 2: The argument above also gives an insight, if not a proof, as to why $\mu = A/B$ in (13) is an integer. Suppose $\mathcal{M}((\mathcal{I}/p^k)^\sim) = \{r_1, r_2, r_3, \dots, r_d\}$. The power of p^k in $A = \prod (m_i - m_j)$ is given by

$$\sum_{i=1}^d \binom{r_i}{2} = \frac{1}{2} \left(\sum_{i=1}^d r_i^2 - \sum_{i=1}^d r_i \right).$$

Now, $\sum_{i=1}^d r_i$ is the cardinality of \mathcal{I} so

$$\sum_{i=1}^d r_i = d.$$

Hence for a set \mathcal{I} which has the minimum power of p^k in A it must be that $\mathcal{M}((\mathcal{I}/p^k)^\sim) = \{r_1, r_2, \dots, r_d\}$ is a solution to

$$\begin{aligned} &\text{minimize } r_1^2 + r_2^2 + \dots + r_d^2 \\ &\text{subject to } r_1 + r_2 + \dots + r_d = d. \end{aligned}$$

On the reals the optimal solution satisfies $r_1 = r_2 = \dots = r_d$. This suggests that the set \mathcal{I} with the smallest power of p^k in A must have roughly an equal number of elements in each congruence class. $\mathcal{I}^* = \{0, 1, 2, \dots, d-1\}$ is one such set. Thus the power of p^k is smaller in $B = \prod (i - j)$ than in $A = \prod (m_i - m_j)$ for each p and k , and, if the reasoning is to be trusted, $\mu = A/B$ is an integer.

To finish the proof of Theorem 4 we will derive the following bounds on $\tilde{\chi}_k$.

Lemma 3: If $\mathcal{I} \subseteq [0 : p^M - 1]$ is a universal sampling set of size d then

$$\left\lfloor \frac{d}{p^k} \right\rfloor \leq \tilde{\chi}_k(s) \leq \left\lceil \frac{d}{p^k} \right\rceil, \quad s \in [0 : p^k - 1], 0 \leq k \leq M. \quad (15)$$

It follows immediately from (15) that if \mathcal{I} is a universal sampling set then

$$|\tilde{\chi}_k(a) - \tilde{\chi}_k(b)| \leq 1, \quad a, b \in [0 : p^k - 1].$$

This is condition (ii), and with this result the proof of Theorem 4 will be complete. Incidentally, for the case $\mathcal{I} = \mathcal{I}^*$, (15) is a simple consequence of (9) and (8).

The argument for Lemma 3 is through constructing submatrices of the Fourier matrix of known rank to obtain upper and lower bounds for $\tilde{\chi}_k$. The first step is to build a particular model submatrix, and this requires some bookkeeping.

Let $\mathcal{I} \subseteq [0 : p^M - 1]$, at this point not assumed to be a universal sampling set. Fix $k \leq M$ and $s \in [0 : p^k - 1]$, and recall that we let

$$\mathcal{I}_{ks} = \{i \in \mathcal{I} : i \equiv s \pmod{p^k}\}.$$

The set \mathcal{I}_{ks} has $\tilde{\chi}_k(s)$ elements. List them, in numerical order, as $i_0, i_1, i_2, \dots, i_c$, where we put $c = \tilde{\chi}_k(s) - 1$ to simplify notation. Let r be a positive integer and define the column vector of length c by

$$\mathfrak{z}^r = [\zeta_N^{i_0 r} \quad \zeta_N^{i_1 r} \quad \zeta_N^{i_2 r} \quad \dots \quad \zeta_N^{i_c r}]^T.$$

Now let \mathfrak{Z}^r be the $c \times p^k$ matrix obtained by repeating p^k copies of the column \mathfrak{z}^r :

$$\mathfrak{Z}^r = \underbrace{[\mathfrak{z}^r \quad \mathfrak{z}^r \quad \mathfrak{z}^r \quad \dots \quad \mathfrak{z}^r]}_{p^k \text{ times}},$$

and let \mathfrak{D}^s be the $p^k \times p^k$ diagonal matrix

$$\mathfrak{D}^s = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 \\ 0 & \zeta_{p^k}^s & 0 & \dots & 0 \\ 0 & 0 & \zeta_{p^k}^{2s} & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \zeta_{p^k}^{(p^k-1)s} \end{bmatrix}.$$

Finally, let $k' = M - k$, and set

$$\mathcal{J}_{k'r} = \{0 \cdot p^{k'} + r, 1 \cdot p^{k'} + r, 2 \cdot p^{k'} + r, \dots, (p^k - 1)p^{k'} + r\}. \quad (16)$$

From the Fourier matrix \mathcal{F} we choose c rows indexed by \mathcal{I}_{ks} and p^k columns indexed by $\mathcal{J}_{k'r}$. The result of these choices, we claim, results in

$$E_{\mathcal{I}_{ks}}^T \mathcal{F} E_{\mathcal{J}_{k'r}} = \mathfrak{Z}^r \mathfrak{D}^s. \quad (17)$$

After the preparations, the derivation of (17) is straightforward. The (a, b) -entry of $E_{\mathcal{I}_{ks}}^\top \mathcal{F} E_{\mathcal{J}_{k'r}}$ is

$$\begin{aligned}\zeta_N^{i_a(bp^{k'}+r)} &= \exp\left(-\frac{(2\pi i)i_a(bp^{M-k}+r)}{p^M}\right) \\ &= \exp\left(-\frac{(2\pi i)i_a b}{p^k}\right) \exp\left(-\frac{(2\pi i)i_a r}{p^M}\right).\end{aligned}$$

But now recall that, by definition, when $i_a \in \mathcal{I}_{ks}$ is divided by p^k it leaves a remainder of s , and thus

$$\begin{aligned}&\exp\left(-\frac{(2\pi i)i_a b}{p^k}\right) \exp\left(-\frac{(2\pi i)i_a r}{p^M}\right) \\ &= \exp\left(-\frac{(2\pi i)s b}{p^k}\right) \exp\left(-\frac{(2\pi i)i_a r}{p^M}\right) \\ &= \zeta_{p^k}^{sb} \zeta_N^{i_a r}.\end{aligned}$$

This construction is the basis for the proof of Lemma 3, but applied in block form.

Proof of Lemma 3: To deduce the upper bound $\tilde{\chi}(s) \leq \lceil d/p^k \rceil$ we begin by letting

$$\mathcal{J} = \mathcal{J}_{k'0} \cup \mathcal{J}_{k'1} \cup \mathcal{J}_{k'2} \cup \dots \cup \mathcal{J}_{k'd'}, \quad d' = \left\lceil \frac{d}{p^k} \right\rceil - 1,$$

where $\mathcal{J}_{k'r}$ is defined as in (16). Note that \mathcal{J} is a union of $\lceil d/p^k \rceil$ disjoint sets. Each $\mathcal{J}_{k'r}$, $0 \leq r \leq d' = \lceil d/p^k \rceil - 1$ indexes the choice of p^k columns from \mathcal{F} and applying (17) we have

$$\begin{aligned}E_{\mathcal{I}_{ks}}^\top \mathcal{F} E_{\mathcal{J}} &= \begin{bmatrix} E_{\mathcal{I}_{ks}}^\top \mathcal{F} E_{\mathcal{J}_{k'0}} & E_{\mathcal{I}_{ks}}^\top \mathcal{F} E_{\mathcal{J}_{k'1}} & \dots & E_{\mathcal{I}_{ks}}^\top \mathcal{F} E_{\mathcal{J}_{k'd'}} \end{bmatrix} \\ &= \begin{bmatrix} \mathfrak{Z}^0 \mathfrak{D}^s & \mathfrak{Z}^1 \mathfrak{D}^s & \dots & \mathfrak{Z}^{d'} \mathfrak{D}^s \end{bmatrix} \\ &= \begin{bmatrix} \mathfrak{Z}^0 & \mathfrak{Z}^1 & \dots & \mathfrak{Z}^{d'} \end{bmatrix} \begin{bmatrix} \mathfrak{D}^s & 0 & \dots & 0 \\ 0 & \mathfrak{D}^s & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \mathfrak{D}^s \end{bmatrix}.\end{aligned}$$

The diagonal matrix in this product is invertible, hence

$$\begin{aligned}\text{Rank of } E_{\mathcal{I}_{ks}}^\top \mathcal{F} E_{\mathcal{J}} &= \text{Rank of } \begin{bmatrix} \mathfrak{Z}^0 & \mathfrak{Z}^1 & \mathfrak{Z}^2 & \dots & \mathfrak{Z}^{d'} \end{bmatrix} \\ &\leq \text{Number of distinct columns} = \left\lceil \frac{d}{p^k} \right\rceil.\end{aligned}\tag{18}$$

Now, the number of columns of $E_{\mathcal{I}_{ks}}^\top \mathcal{F} E_{\mathcal{J}}$ is equal to

$$\begin{aligned}|\mathcal{J}| &= |\mathcal{J}_{k'0} \cup \mathcal{J}_{k'1} \cup \mathcal{J}_{k'2} \cup \dots \cup \mathcal{J}_{k'd'}| \\ &= \sum_{r=0}^{\lceil d/p^k \rceil - 1} |\mathcal{J}_{k'r}| = p^k \lceil d/p^k \rceil \geq d,\end{aligned}$$

so there are at least d columns. Hence if \mathcal{I} is a universal sampling set of size d then $E_{\mathcal{I}}^\top \mathcal{F} E_{\mathcal{J}}$ must be of full row rank. In particular, since $\mathcal{I}_{ks} \subseteq \mathcal{I}$, it must be that $E_{\mathcal{I}_{ks}}^\top \mathcal{F} E_{\mathcal{J}}$ is also of full row rank, for each s . Next, the number of rows in $E_{\mathcal{I}_{ks}}^\top \mathcal{F} E_{\mathcal{J}}$ is equal to $|\mathcal{I}_{ks}| = \tilde{\chi}_k(s)$ by definition. From (18) we know that the rank of $E_{\mathcal{I}_{ks}}^\top \mathcal{F} E_{\mathcal{J}}$ is at most $\lceil d/p^k \rceil$, and so we have

$$(\text{Number of rows}) \tilde{\chi}_k(s) \leq \left\lceil \frac{d}{p^k} \right\rceil.$$

The proof of the lower bound $\tilde{\chi}_k(s) \geq \lfloor d/p^k \rfloor$ is very similar. This time we construct a set \mathcal{J} with $|\mathcal{J}| \leq d$, and observe that if \mathcal{I} is a universal sampling set of size d , then $E_{\mathcal{I}}^\top \mathcal{F} E_{\mathcal{J}}$ is of full column rank.

Let

$$\mathcal{J} = \mathcal{J}_{k'0} \cup \mathcal{J}_{k'1} \cup \mathcal{J}_{k'2} \cup \dots \cup \mathcal{J}_{k'd''}, \quad d'' = \lfloor d/p^k \rfloor - 1.$$

Then just as above,

$$\begin{aligned}\text{Rank of } E_{\mathcal{I}_{ks}}^\top \mathcal{F} E_{\mathcal{J}} &= \text{Rank of } \begin{bmatrix} \mathfrak{Z}^0 & \mathfrak{Z}^1 & \mathfrak{Z}^2 & \dots & \mathfrak{Z}^{d''} \end{bmatrix} \\ &\leq \text{Number of distinct columns} = \left\lfloor \frac{d}{p^k} \right\rfloor.\end{aligned}$$

The number of rows of $E_{\mathcal{I}_{ks}}^\top \mathcal{F} E_{\mathcal{J}}$ is $|\mathcal{I}_{ks}| = \tilde{\chi}_k(s)$, and so we must have

$$\text{Rank of } E_{\mathcal{I}_{ks}}^\top \mathcal{F} E_{\mathcal{J}} \leq \min\{\lfloor d/p^k \rfloor, \tilde{\chi}_k(s)\}.\tag{19}$$

Furthermore,

$$E_{\mathcal{I}}^\top \mathcal{F} E_{\mathcal{J}} = \begin{bmatrix} E_{\mathcal{I}_{k0}}^\top \mathcal{F} E_{\mathcal{J}} \\ E_{\mathcal{I}_{k1}}^\top \mathcal{F} E_{\mathcal{J}} \\ \vdots \\ E_{\mathcal{I}_{k(p^k-1)}}^\top \mathcal{F} E_{\mathcal{J}} \end{bmatrix},$$

whence

$$\begin{aligned}\text{Row rank of } E_{\mathcal{I}}^\top \mathcal{F} E_{\mathcal{J}} &\leq \text{Rank of } E_{\mathcal{I}_{k0}}^\top \mathcal{F} E_{\mathcal{J}} + \text{Rank of } E_{\mathcal{I}_{k1}}^\top \mathcal{F} E_{\mathcal{J}} \\ &\quad + \dots + \text{Rank of } E_{\mathcal{I}_{k(p^k-1)}}^\top \mathcal{F} E_{\mathcal{J}} \\ &\leq \sum_{s=0}^{p^k-1} \min\{\lfloor d/p^k \rfloor, \tilde{\chi}_k(s)\}.\end{aligned}\tag{20}$$

Now, the number of columns indexed by \mathcal{J} is $p^k \lfloor d/p^k \rfloor \leq d$. Hence if \mathcal{I} is a universal sampling set of size d , we need $E_{\mathcal{I}}^\top \mathcal{F} E_{\mathcal{J}}$ to be of full column rank. From (20), this means we must have

$$(\text{Number of columns}) p^k \lfloor d/p^k \rfloor \leq \sum_{s=0}^{p^k-1} \min\{\lfloor d/p^k \rfloor, \tilde{\chi}_k(s)\}.$$

This inequality will not be satisfied unless $\lfloor d/p^k \rfloor \leq \tilde{\chi}_k(s)$ for all s . This completes the proof. ■

Remark 3: For many values of d , it is enough to prove one side of the inequality (15). If we know that $\tilde{\chi}_k(s) \leq \lceil d/p^k \rceil$, then from $\sum_s \tilde{\chi}_k(s) = d$ and a recurrence relation (23), below, it is possible to prove that $\lfloor d/p^k \rfloor \leq \tilde{\chi}_k(s)$. Such cases include

- 1) $N = p^M$, $d = c_0 p^k + c_1 p^{k-1}$ for $c_0, c_1 \in \{0, 1, 2, \dots, p-1\}$.
- 2) $N = 2^M$, $d = c_0 2^k + c_1 2^{k-1} + c_2 2^{k-2}$ for $c_0, c_1, c_2 \in \{0, 1\}$.
- 3) $N = 2^M$, $d = 2^k + 2^{k-1} + 2^{k-2} + \dots + 2^{k-r+1}$ for some r ,

C. Digit Reversal and Universal Sampling Sets

There is another interesting characterization of universal sampling sets in terms of *digit reversal*. Expanding in base p , any integer $a \in [0 : p^m - 1]$, $m \geq 1$, can be written uniquely as

$$a = \alpha_0 + \alpha_1 p + \alpha_2 p^2 + \cdots + \alpha_{m-1} p^{m-1},$$

where the α 's are in $[0 : p - 1]$. We define a permutation $\pi_m : [0 : p^m - 1] \rightarrow [0 : p^m - 1]$ by

$$\begin{aligned} \pi_m(\alpha_0 + \alpha_1 p + \alpha_2 p^2 + \cdots + \alpha_{m-1} p^{m-1}) = \\ \alpha_{m-1} + \alpha_{m-2} p + \alpha_{m-3} p^2 + \cdots + \alpha_0 p^{m-1}. \end{aligned}$$

The α 's are the digits in the base p expansion of $a \in [0 : p^m - 1]$ and applying π_m to a produces the number in $[0 : p^m - 1]$ with the digits reversed. For example (an example we will use again in Section V), take $[0 : 7]$. Then $\pi_3([0 : 7]) = \{0, 4, 2, 6, 1, 5, 3, 7\}$ in that order. Such digit reversing permutations were used in [10] to find rank-one submatrices of the Fourier matrix.

The issue for universal sampling sets is how the numbers $\pi_M(\mathcal{I})$ are dispersed within the interval $[0 : p^M - 1]$, where, as before, $N = p^M$. To make this precise, take $k \geq 1$ and partition $[0 : p^M - 1]$ into p^k equal parts:

$$[0 : p^M - 1] = \bigcup_{a=0}^{p^k-1} [ap^{k'} : (a+1)p^{k'} - 1], \quad k' = M - k.$$

For any $\mathcal{J} \subseteq [0 : p^M - 1]$ and $a \in [0 : p^k - 1]$, let

$$\phi_k(a; \mathcal{J}) = |\mathcal{J} \cap [ap^{k'}, (a+1)p^{k'} - 1]|.$$

We say that \mathcal{J} is *uniformly dispersed* in $[0, p^M - 1]$ if

$$|\phi_k(a; \mathcal{J}) - \phi_k(b; \mathcal{J})| \leq 1 \quad (21)$$

for all $a, b \in [0 : p^k - 1]$, and $1 \leq k \leq M$. Thus \mathcal{J} is uniformly dispersed if roughly equal numbers of its elements are in each of the intervals $[ap^{k'} : (a+1)p^{k'} - 1]$ for all $1 \leq k \leq M$, $k' = M - k$.

We will show

$$\phi_k(\pi_k(a); \pi_M(\mathcal{I})) = \tilde{\chi}_k(a), \quad a \in [0 : p^k - 1]. \quad (22)$$

Thus, to the three equivalent conditions in Theorem 4 we can add a fourth:

(iv) $\pi_M(\mathcal{I})$ is uniformly dispersed.

The derivation of (22) uses the following lemma.

Lemma 4: If $j \in [0 : p^M - 1]$ is given by $j = b + ap^{k'}$, $0 \leq b \leq p^{k'} - 1$, then $\pi_M(j) = \pi_k(a) + p^k \pi_{k'}(b)$.

The proof is straightforward, and the argument for (22) then goes very quickly. As defined, for any index set \mathcal{J} , $\phi_k(a; \mathcal{J})$ is the number of elements in \mathcal{J} that lie in $[ap^{k'} : (a+1)p^{k'} - 1]$, and these are precisely the $j \in \mathcal{J}$ of the form $ap^{k'} + b$ with

$0 \leq b \leq p^{k'} - 1$. Thus for $i \in [0 : p^k - 1]$,

$$\begin{aligned} \phi_k(\pi_k(i); \pi_M(\mathcal{I})) &= \\ &\text{the number of } j \in \pi_M(\mathcal{I}) \\ &\text{of the form } \pi_k(i)p^{k'} + b, \quad 0 \leq b \leq p^{k'} - 1 \\ &= \text{number of } j \in \mathcal{I} \text{ of the form} \\ &\quad p^k \pi_{k'}(b) + i, \quad 0 \leq b \leq p^{k'} - 1 \text{ (from Lemma 4)} \\ &= \text{number of } j \in \mathcal{I} \text{ that leave a remainder of} \\ &\quad i \text{ on dividing by } p^k \\ &= \tilde{\chi}_k(i). \end{aligned}$$

V. STRUCTURE AND ENUMERATION OF UNIVERSAL SAMPLING SETS

In this section we analyze in detail the structure of universal sampling sets. Specifically we show that when $N = p^M$ is a prime power such a set \mathcal{I} is the disjoint union of smaller, *elementary universal sets* that depend on the base p expansion of $|\mathcal{I}|$. The method is algorithmic, allowing us to construct universal sets of a given size, and to find a formula that counts the number of universal sets as a function of p^M and $|\mathcal{I}|$. In particular the formula answers the question: How likely is it that a randomly chosen index set is universal? Not very likely, but there are several subtle aspects to the answer. For example, we exhibit plots of the counting function showing some striking phenomena depending on the prime p . Our approach is via *maximal universal sampling sets* which, in turn, enter naturally in studying the relationship between universal sampling sets and uncertainty principles. We take up the latter topic in the next section.

A. A Recurrence Relation and Tree for $\tilde{\chi}$

When $N = p^M$ the condition that an index set be a universal sampling set depends on the values of $\tilde{\chi}_k$ for different k . To study this we use a recurrence relation in k for $\tilde{\chi}_k(a)$. The formula holds even when N is not a prime power.

Lemma 5: Let $\mathcal{I} \subseteq [0 : N - 1]$. Then

$$\tilde{\chi}_{k-1}(a) = \sum_{j=0}^{p-1} \tilde{\chi}_k(a + jp^{k-1}), \quad (23)$$

for all $a \in [0 : p^{k-1} - 1]$.

Proof: An integer $x \in \mathcal{I}$ that leaves a remainder of a when divided by p^{k-1} is of the form $x = \alpha p^{k-1} + a$. Let $\alpha = \beta p + \gamma$ for $\gamma \in [0 : p - 1]$. Then $x = \beta p^k + \gamma p^{k-1} + a$, that is, x leaves a remainder of either $0 \cdot p^{k-1} + a, 1 \cdot p^{k-1} + a, 2 \cdot p^{k-1} + a, \dots$ or $(p-1) \cdot p^{k-1} + a$ on dividing by p^k . The result follows. ■

When $N = p^M$ the recurrence formula and the relation it expresses between conjugacy classes has an appealing interpretation in terms of a p -ary tree. Several arguments in this section will be based on this configuration.

Let $\mathcal{I} \subseteq [0 : p^M - 1]$. We construct a tree with $M+1$ levels and p^k nodes in level k , $0 \leq k \leq M$. The nodes in level k are identified by a pair (k, a) , with $a \in [0 : p^{k-1}]$. Call the nodes at the level M the leaves. At the node (k, a) we imagine placing the congruence class $\mathcal{I}_{ka} = \{i \in \mathcal{I} : i \equiv a \pmod{p^k}\}$.

The root is $\mathcal{I}_{00} = \mathcal{I}$ and the nodes at the leaves host the sets \mathcal{I}_{Ma} , $a \in [0 : p^M - 1]$, each of which is either a singleton or empty. We assign a weight of $\tilde{\chi}_k(a) = |\mathcal{I}_{ka}|$ to the node (k, a) . Further, at each level we arrange the nodes according to the digit reversing permutation, i.e., nodes at level k are arranged as $\pi_k([0 : p^k - 1])$, where π_k is the digit reversing permutation from Section IV-C. (This is similar to the starting step of the FFT algorithm, where the indices are sorted according to the reversed digits.) Figure 2 shows the case $N = 2^3$, a binary tree with four levels, $k = 0, 1, 2, 3$. In the third level of the tree the nodes are ordered 0, 4, 2, 6, 1, 5, 3, 7, which is $\pi_3([0 : 7])$. Then:

- 1) The set \mathcal{I}_{ka} at level k is the disjoint union of the sets at its children nodes at level $k + 1$.
- 2) The value of $\tilde{\chi}_k(a)$ at the node (k, a) is the sum of the values of $\tilde{\chi}_{k+1}$ at its children nodes at level $k + 1$. In other words, the weight of a parent is the sum of the weights of its children; this is the recurrence relation. Consequently, the value of $\tilde{\chi}_k$ at any node is the sum of the values of $\tilde{\chi}_M$ at the leaves at level M descended from the node.

For example, in Figure 2 we have

$$\begin{aligned}\tilde{\chi}_0(0) &= \sum_{a=0}^7 \tilde{\chi}_3(a), \\ \tilde{\chi}_1(0) &= \sum_{a=0}^3 \tilde{\chi}_3(2a), \\ \tilde{\chi}_1(1) &= \sum_{a=0}^3 \tilde{\chi}_3(2a+1),\end{aligned}$$

and so on.

In fact, a more general conclusion is the following: Fix a level k . Then the value of $\tilde{\chi}_r$ at any node (r, a) , for $r \leq k$ is the sum of the values of $\tilde{\chi}_k$ at the level- k nodes descending from the tree node (r, a) .

When the root is $[0 : p^M - 1]$, the extreme case, the leaves are all singletons and the nodes at level k are each of weight p^{M-k} .

B. Elementary and Maximal Sets

To study the structure of universal sampling sets we need a series of definitions. When N is a prime power the building blocks are the elementary sets:

Definition 4: A set $\mathcal{E} \subseteq [0 : p^M - 1]$ is a k -elementary set if

$$\tilde{\chi}_k(a) = 1, \quad \text{for all } a \in [0 : p^k - 1].$$

Note that $|\mathcal{E}| = p^k$.

As a first application of the formula (23) we can add the adjective “universal” to the description of elementary sets.

Lemma 6: A k -elementary set \mathcal{E} is a universal sampling set.

Proof: From $\tilde{\chi}_k(a) = 1$ and (23) it follows that \mathcal{E} has an equal number of elements in each congruence class modulo p^s , $s \leq k$. More precisely,

$$\tilde{\chi}_s(a) = p^{k-s}, \quad (24)$$

for all $s \leq k$. Also from (23), for $s > k$ all the congruence classes are of size 0 or 1, i.e.

$$\tilde{\chi}_s(a) \in \{0, 1\}. \quad (25)$$

Therefore

$$|\tilde{\chi}_s(a) - \tilde{\chi}_s(b)| \leq 1,$$

for all $a, b \in [0 : p^k - 1]$ and all s , and we conclude that \mathcal{E} is a universal sampling set. ■

Next, a fruitful approach to understanding the structure of universal sampling sets is to ask how well an arbitrary index set is approximated from within by universal sets.

Definition 5: Let $\mathcal{I} \subseteq [0 : N - 1]$. A *maximal universal sampling set* for \mathcal{I} is a universal sampling set of largest cardinality that is contained in \mathcal{I} .

Note that the definition does not require N to be a prime power, though this will most often be the case. There is an allied notion of a minimal universal set. We define this in Subsection V-E below, and show how they are related to maximal sets. Maximal and minimal sets enter naturally and together in connection with uncertainty principles, discussed in Section VI.

Finding a maximal universal sampling set for a given \mathcal{I} is a finitary process, so existence is not an issue. However, maximal universal sampling sets need not be unique. For example, take $N = 3^2$ and $\mathcal{I} = \{0, 1, 2, 3, 6\}$. The set \mathcal{I} is not itself a universal sampling set, and both $\{0, 1, 2, 3\}$ and $\{0, 1, 2, 6\}$ are maximal universal sampling sets contained in \mathcal{I} .

Despite the lack of uniqueness it will be convenient to have a notation, and we let $\Omega(\mathcal{I})$ denote a generic maximal universal sampling set in \mathcal{I} . The cardinality $|\Omega(\mathcal{I})|$ is well-defined; by definition $|\mathcal{J}| \leq |\Omega(\mathcal{I})|$ for any universal sampling set $\mathcal{J} \subseteq \mathcal{I}$.

Elementary sets and maximal sets are related through an important construction of an elementary set.

Definition 6: Let $\mathcal{I} \subseteq [0 : p^M - 1]$ and let \bar{k} be the largest integer such that no congruence class in $\mathcal{I}/p^{\bar{k}}$ is empty. (It might be that $\bar{k} = 0$.) Let \mathcal{I}_k^\dagger denote an elementary set obtained by choosing one element from each congruence class in $\mathcal{I}/p^{\bar{k}}$.

By Lemma 6, \mathcal{I}_k^\dagger is a universal sampling set, and is of order $p^{\bar{k}}$. We now have

Theorem 5: Let $\mathcal{I} \subseteq [0 : p^M - 1]$, and \mathcal{I}_k^\dagger as above. Then

- (i) $p^{\bar{k}} \leq |\Omega(\mathcal{I})| < p^{\bar{k}+1}$.
- (ii) There exists a maximal universal sampling set contained in \mathcal{I} and containing \mathcal{I}_k^\dagger .

Proof: The lower bound in (i) follows from the definition of a maximal set and the comments above,

$$p^{\bar{k}} = |\mathcal{I}_k^\dagger| \leq |\Omega(\mathcal{I})|.$$

To prove the upper bound, suppose $\mathcal{J} \subseteq \mathcal{I}$ has $|\mathcal{J}| \geq p^{\bar{k}+1}$. By the definition of \bar{k} at least one congruence class in $\mathcal{J}/p^{\bar{k}+1}$ is empty, so $\tilde{\chi}_{\bar{k}+1}(a; \mathcal{J}) = 0$ for some $a \in [0 : p^{\bar{k}+1} - 1]$. From the cardinality equation (6),

$$\sum_{\ell=0}^{p^{\bar{k}+1}-1} \tilde{\chi}_{\bar{k}+1}(\ell; \mathcal{J}) = |\mathcal{J}| \geq p^{\bar{k}+1}.$$




Fig. 2: A tree representing the relations between the congruence classes, and the recurrence relation satisfied by $\tilde{\chi}_k(a)$. The value of $\tilde{\chi}_k(a)$ at any node is the sum of the values of $\tilde{\chi}_k(a)$ at its children nodes in level $k + 1$.

This implies that at least one congruence class in $\mathcal{I}/p^{\bar{k}+1}$ has at least two elements, or $\tilde{\chi}_{\bar{k}+1}(b; \mathcal{I}) \geq 2$ for some b . We then have

$$|\tilde{\chi}_{\bar{k}+1}(b; \mathcal{I}) - \tilde{\chi}_{\bar{k}+1}(a; \mathcal{I})| = 2 > 1,$$

and \mathcal{I} cannot be a universal sampling set.

For part (ii), we first show that any maximal universal sampling set $\Omega(\mathcal{I})$ set must contain at least one element from each congruence class in $\mathcal{I}/p^{\bar{k}}$. By way of contradiction, suppose that $\tilde{\chi}_{\bar{k}}(a; \Omega(\mathcal{I})) = 0$ for some a . Since $\Omega(\mathcal{I})$ is universal we must then have $\tilde{\chi}_{\bar{k}}(b; \Omega(\mathcal{I})) \leq 1$ for all b . By (6),

$$|\Omega(\mathcal{I})| = \sum_{b=0}^{p^{\bar{k}}-1} \tilde{\chi}_{\bar{k}}(b; \Omega(\mathcal{I})) < p^{\bar{k}},$$

contradicting the lower bound in (i).

Let $\mathcal{K} \subseteq \Omega(\mathcal{I})$ be an elementary set, of size $p^{\bar{k}}$, that contains one element from each congruence class in $\mathcal{I}/p^{\bar{k}}$, guaranteed to exist from what we just showed. Assuming $\mathcal{K} \neq \mathcal{I}_k^\dagger$, since otherwise we are done, we will use \mathcal{K} and $\Omega(\mathcal{I})$ to construct a (new) maximal universal set that contains \mathcal{I}_k^\dagger .

Set up a p -ary tree, as above, with root $\mathcal{I}_{00} = \Omega(\mathcal{I})$ and (ℓ, a) -node the congruence class

$$\Omega_{\ell a} = \{i \in \Omega(\mathcal{I}) : i \equiv a \pmod{p^\ell}\}, \quad |\Omega_{\ell a}| = \tilde{\chi}_\ell(a),$$

for $a \in [0 : p^\ell - 1]$. Recall that $\Omega_{\ell a}$, at level ℓ , is the disjoint union of the sets at its children nodes at level $\ell + 1$.

Figure 3 is an example for $p = 3$ and $M \geq 3$, showing only three levels for reasons of space. The shading has to do with the rest of the proof, as we now explain.

Both \mathcal{I}_k^\dagger and \mathcal{K} are assembled by choosing single elements from sets at the nodes in the \bar{k} -level (call these the assembly nodes) for a total of $p^{\bar{k}}$ elements for \mathcal{I}_k^\dagger and \mathcal{K} each. Observe that the sets at the nodes in the $\bar{k} + 1$ level are either empty or singletons. This is so because by definition of \bar{k} there must be some $a \in [0 : p^{\bar{k}+1} - 1]$ for which $\tilde{\chi}_{\bar{k}+1}(a) = 0$, and hence by universality $\tilde{\chi}_{\bar{k}+1}(b) \leq 1$ for all $b \in [0 : p^{\bar{k}+1} - 1]$. And

then, according to how the tree is structured, the sets at all nodes farther down in the tree must as well be either empty or singletons.

Let $\mathcal{L} \supseteq \mathcal{I}_k^\dagger$ be the set of elements in \mathcal{I} that leave the same remainders as do the elements in \mathcal{I}_k^\dagger when divided by $p^{\bar{k}+1}$, more precisely,

$$\mathcal{L} = \{j \in \mathcal{I} : j \equiv i \pmod{p^{\bar{k}+1}} \text{ for some } i \in \mathcal{I}_k^\dagger\}.$$

Likewise let $\mathcal{L}' \supseteq \mathcal{K}$ be

$$\mathcal{L}' = \{j \in \mathcal{I} : j \equiv i \pmod{p^{\bar{k}+1}} \text{ for some } i \in \mathcal{K}\}.$$

\mathcal{L} is the union of the assembly nodes for \mathcal{I}_k^\dagger and \mathcal{L}' is the union of the assembly nodes for \mathcal{K} . The collections may overlap.

We color a node red if it contributes to \mathcal{L} and blue if it contributes to \mathcal{L}' , and both red and blue (otherwise known as purple) if it contributes to both \mathcal{L} and \mathcal{L}' . In the figure we take $\bar{k} = 1$, so \mathcal{I}_k^\dagger and \mathcal{K} live at the middle level in the tree, as shown.

Focus on each red node in turn. The red node contains an element in \mathcal{I}_k^\dagger , say i .

- 1) If $\Omega(\mathcal{I})$ contains an element from this red node, say j (which may or may not be equal to i), we replace $j \in \Omega(\mathcal{I})$ with i . This neither changes the size of $\Omega(\mathcal{I})$ nor the universality.
- 2) Now suppose $\Omega(\mathcal{I})$ does not contain an element from this red node. We know that the sibling blue node (i.e. the blue node that shares the parent with this red node) contains an element of \mathcal{K} (and hence of $\Omega(\mathcal{I})$), say j . Replace $j \in \Omega(\mathcal{I})$ with i . This neither changes the size, nor the universality; we are just exchanging one element from a node with its sibling, so the value of $\tilde{\chi}$ at the parent node does not change.

These operations preserve size and universality, and repeating them for each red node ensures that the resultant set contains \mathcal{I}_k^\dagger .

■

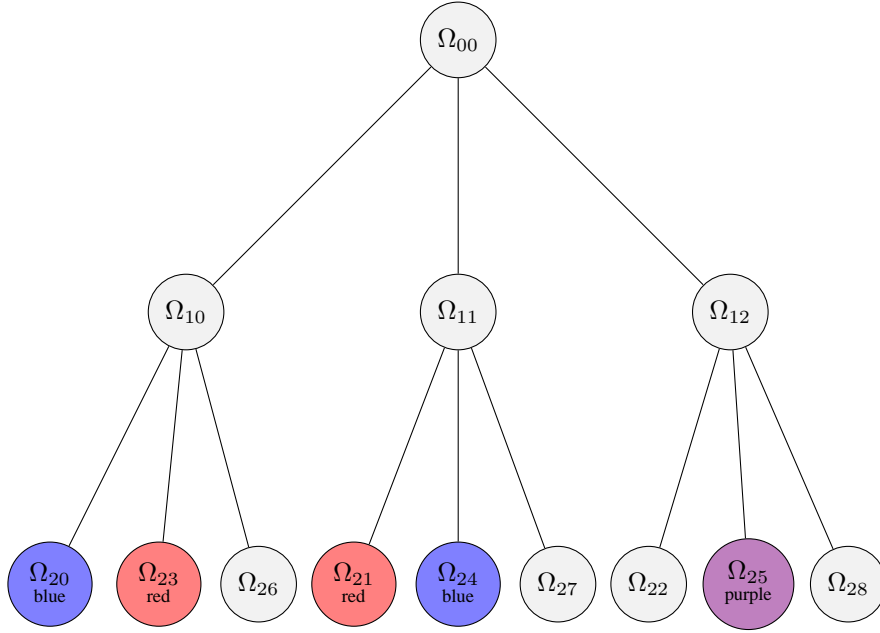


Fig. 3: The congruence class tree for $\Omega(\mathcal{I})$. The node $\Omega_{\ell a}$ is the congruence class of a modulo p^ℓ in $\Omega(\mathcal{I})$, so that $\Omega_{00} = \Omega(\mathcal{I})$ and $|\Omega_{\ell a}| = \tilde{\chi}_\ell(a)$.

A stronger version of the upper bound in (i) is the following.

Corollary 1: Let $\mathcal{I} \subseteq [0 : p^M - 1]$ and let \underline{k} be the smallest integer such that $\tilde{\chi}_{\underline{k}}(a; \mathcal{I}) = 0$ for some a . Then,

$$|\Omega(\mathcal{I})| \leq |\{a : \tilde{\chi}_{\underline{k}}(a; \mathcal{I}) \neq 0\}|.$$

Proof: From the definition of \underline{k} , we have $\tilde{\chi}_{\underline{k}}(a_0; \mathcal{I}) = 0$ for some a_0 . Hence by universality, $\Omega(\mathcal{I})$ must satisfy $|\tilde{\chi}_{\underline{k}}(b; \Omega(\mathcal{I}))| \leq 1$ for all b , an observation we used above and will use again. From the cardinality equation (6)

$$|\Omega(\mathcal{I})| = \sum_b \tilde{\chi}_{\underline{k}}(b; \Omega(\mathcal{I})) \leq |\{a : \tilde{\chi}_{\underline{k}}(a; \mathcal{I}) \neq 0\}|.$$

Ultimately we will show that when $N = p^M$ any maximal universal sampling set, and in particular any universal sampling set, is a disjoint union of elementary sets. In general, however, the union of two disjoint, elementary sets need not be universal. For example, take $N = 2^3$, $\mathcal{E} = \{0, 1\}$, $\mathcal{E}' = \{4, 5\}$. Then \mathcal{E} and \mathcal{E}' are elementary but their union $\mathcal{E} \cup \mathcal{E}' = \{0, 1, 4, 5\}$ is not universal. What is needed is a kind of independence condition on a collection of elementary sets. The following lemma, whose converse we will also show, makes this latter point precise and introduces the main features of the structure of universal sets.

Lemma 7: Let $N = p^M$. Suppose there exists a finite sequence of nonincreasing integers $k_1 \geq k_2 \geq \dots \geq 0$ and sets $\mathcal{E}_r \subseteq [0 : N - 1]$, $r = 1, 2, \dots$, such that

- (i) \mathcal{E}_r is k_r -elementary.
- (ii) For each $r \geq 1$

$$\mathcal{E}_r \cap \left(\bigcup_{j=1}^{r-1} \mathcal{L}_j \right) = \emptyset,$$

where

$$\mathcal{L}_j = \{x \in [0 : N - 1] : x \equiv e \pmod{p^{k_j+1}} \text{ for some } e \in \mathcal{E}_j\}.$$

Let

$$\mathcal{I} = \bigcup_r \mathcal{E}_r.$$

Then \mathcal{I} is a universal sampling set.

Obviously it is condition (ii) that requires further comment. The set \mathcal{L}_r is defined much as in the proof of Theorem 5, and we will illustrate the point of (ii) again by means of a tree. Observe first that the \mathcal{E}_r are disjoint. This follows from (ii), since $\mathcal{L}_j \supseteq \mathcal{E}_j$.

We build a congruence tree with root the full interval $[0 : N - 1]$. Write this as \mathcal{N}_{00} and write \mathcal{N}_{ka} for the congruence class of a modulo p^k in $[0 : N - 1]$, so that $|\mathcal{N}_{ka}| = \tilde{\chi}_k(a; [0 : N - 1])$. All the nodes represent non-singletons, except the bottom-most level, M . As before, Figure 4 has $p = 3$, $M \geq 3$ and shows the tree only up to the third level.

Suppose $k_1 = 1$, so \mathcal{E}_1 , as an elementary set, contains one element from each node at the middle level in the figure. In turn, suppose \mathcal{E}_1 comes from picking one element from each of the red nodes. The set \mathcal{L}_1 is the union of the red nodes. Now, the set \mathcal{E}_2 comes from choosing one element from each node at the k_2 -level, and the sequence k_r is nonincreasing so \mathcal{E}_2 is drawn from nodes in a level at or higher up in the tree than \mathcal{E}_1 (in this example k_2 is either 1 or 0). Condition (ii) requires that \mathcal{E}_2 be disjoint from the red nodes, not just from \mathcal{E}_1 which is a (small) subset of the red nodes.

In the general case, think of k_1 as large (eventually it will be chosen as in Theorem 5), so \mathcal{E}_1 comes from a level far down the tree from the root, and then $\mathcal{E}_2, \mathcal{E}_3, \dots$ are, at least, no further down since $k_1 \geq k_2 \geq \dots$. Condition (ii) requires that \mathcal{E}_r be assembled from nodes that were not used in assembling

any of the \mathcal{E}_s for $s < r$. It is this property that we exploit to show that $\bigcup_r \mathcal{E}_r$ is universal.

Proof of Lemma 7: Fix r and s with $k_r < s$, and note that

$$\tilde{\chi}_s(a; \mathcal{E}_r) \in \{0, 1\}, \quad a \in [0 : p^s - 1], \quad (26)$$

from (25). Now suppose $\tilde{\chi}_s(a; \mathcal{E}_r) = 1$, so one element in \mathcal{E}_r leaves a remainder of a on dividing by p^s . Then

$$\tilde{\chi}_s(a; \mathcal{E}_t) = 0 \quad \text{for all } t > r, \quad (27)$$

i.e., none of the \mathcal{E}_t for $t > r$ will have an element from the congruence class of a modulo p^s . This follows (just as described for the tree) from $\mathcal{E}_t \cap \mathcal{L}_r = \emptyset$, and also

$$\begin{aligned} \mathcal{L}_r &= \{x \in [0 : N - 1] : x \equiv e \pmod{p^{k_r+1}} \text{ for some } e \in \mathcal{E}_r\} \\ &\supseteq \{x \in [0 : N - 1] : x \equiv e \pmod{p^s} \text{ for some } e \in \mathcal{E}_r\}. \end{aligned}$$

From (26) and (27) we conclude that

$$\sum_r \tilde{\chi}_s(a; \mathcal{E}_r) \in \{0, 1\} \quad (28)$$

for all a , where the sum is over all r with $k_r < s$.

With this we can show that $\mathcal{I} = \bigcup_r \mathcal{E}_r$ is universal. For any s , and for any $a, b \in [0 : p^s - 1]$,

$$\begin{aligned} \tilde{\chi}_s(a; \mathcal{I}) - \tilde{\chi}_s(b; \mathcal{I}) &= \sum_r \tilde{\chi}_s(a; \mathcal{E}_r) - \sum_r \tilde{\chi}_s(b; \mathcal{E}_r) \\ &= \left(\sum_{r(k_r \geq s)} \tilde{\chi}_s(a; \mathcal{E}_r) - \sum_{r(k_r \geq s)} \tilde{\chi}_s(b; \mathcal{E}_r) \right) + \quad (29) \\ &\quad \left(\sum_{r(k_r < s)} \tilde{\chi}_s(a; \mathcal{E}_r) - \sum_{r(k_r < s)} \tilde{\chi}_s(b; \mathcal{E}_r) \right) \\ &= \sum_{r(k_r < s)} \tilde{\chi}_s(a; \mathcal{E}_r) - \sum_{r(k_r < s)} \tilde{\chi}_s(b; \mathcal{E}_r) \\ &\quad \text{(the first two sums cancel, by (24)).} \end{aligned}$$

From (28) we have that $|\tilde{\chi}_s(a; \mathcal{I}) - \tilde{\chi}_s(b; \mathcal{I})| \leq 1$, so \mathcal{I} is universal. ■

C. An Algorithm to Construct Maximal Universal Sets

Consider now the problem of finding a maximal universal sampling set contained in a given $\mathcal{I} \subseteq [0 : p^M - 1]$. Build the congruence class tree with root \mathcal{I} , as in Figure 2, up to level M . The leaves having weight 1 are singletons in \mathcal{I} , and $\tilde{\chi}_k(a; \mathcal{I})$, $a \in [0 : p^k - 1]$, is the total weight at node (k, a) . The problem of constructing $\Omega(\mathcal{I})$ is to pick a subset of the leaves so that the tree with root $\Omega(\mathcal{I})$ is well balanced at each level. By ‘well balanced’ we mean that at any given level, all the subtrees have roughly equal weight, corresponding to the condition $|\tilde{\chi}_k(a; \Omega(\mathcal{I})) - \tilde{\chi}_k(b; \Omega(\mathcal{I}))| \leq 1$. The following algorithm realizes this and provides the value of $|\Omega(\mathcal{I})|$. It marries the construction of elementary sets in Theorem 5 with an iterative version of the method used in the proof of Lemma 7.

Let $\mathcal{I} \subseteq [0 : p^M - 1]$. Initialize with $\mathcal{I}_1 = \mathcal{I}$, and $r = 1$.

- 1) Let k_r be the largest integer such that no congruence class in \mathcal{I}_r/p^{k_r} is empty.

- 2) Construct an elementary set $\mathcal{I}_r^\dagger \subseteq \mathcal{I}_r$ by choosing one element of \mathcal{I}_r from each congruence class modulo p^{k_r} . (There may not be a unique choice, and this is the reason why there may be many universal sets contained in \mathcal{I} .)
- 3) Define $\mathcal{L}_r \supseteq \mathcal{I}_r^\dagger$ by

$$\mathcal{L}_r = \{j \in \mathcal{I} : j \equiv i \pmod{p^{k_r+1}} \text{ for some } i \in \mathcal{I}_r^\dagger\}.$$

- 4) Let $\mathcal{I}_{r+1} = \mathcal{I}_r \setminus \mathcal{L}_r$. Stop if $\mathcal{I}_{r+1} = \emptyset$. Else increment r to $r + 1$ and go to (1).

Note the following:

- (i) At each step of the algorithm the size of \mathcal{I}_r is reduced by $|\mathcal{L}_r| \geq |\mathcal{I}_r^\dagger| = p^{k_r} \geq 1$:

$$|\mathcal{I}_{r+1}| \leq |\mathcal{I}_r| - p^{k_r}.$$

Since $\mathcal{I} = \mathcal{I}_1$ is a finite set, the algorithm terminates at some point.

- (ii) The k_r are nonincreasing:

$$k_1 \geq k_2 \geq k_3 \geq \dots$$

We can now state

Theorem 6: With k_r , $r \geq 1$, defined as above, we have

$$|\Omega(\mathcal{I})| = \sum_r p^{k_r}. \quad (30)$$

One possible maximal universal sampling set is

$$\Omega(\mathcal{I}) = \bigcup_r \mathcal{I}_r^\dagger. \quad (31)$$

By construction this is a disjoint union.

Here is an example of the algorithm in action. Let $N = 2^5$ and

$$\mathcal{I} = \{0, 1, 2, 3, 4, 6, 7, 8, 9, 10, 12, 14, 15\} = \mathcal{I}_1.$$

- 1) ($r = 1$) Note that $\tilde{\chi}_3(5; \mathcal{I}_1) = 0$, and that no values $\tilde{\chi}_2(i; \mathcal{I}_1)$ are zero. Hence $k_1 = 2$. Form \mathcal{I}_1^\dagger by taking one element from each congruence class in \mathcal{I}_1 modulo $2^{k_1} = 4$, e.g. $\mathcal{I}_1^\dagger = \{0, 1, 2, 3\}$. Then $\mathcal{L}_1 = \{0, 1, 2, 3, 8, 9, 10\}$ is the set of all elements of \mathcal{I}_1 that leave a remainder of 0, 1, 2 or 3 on dividing by $2^{k_1+1} = 8$. Removing such numbers from \mathcal{I}_1 , we have $\mathcal{I}_2 = \mathcal{I}_1 \setminus \mathcal{L}_1 = \{4, 6, 7, 12, 14, 15\}$.
- 2) ($r = 2$) Now $\tilde{\chi}_2(1; \mathcal{I}_2) = 0$ while $\tilde{\chi}_1(0; \mathcal{I}_2)$, $\tilde{\chi}_1(1; \mathcal{I}_2) \neq 0$ so $k_2 = 1$. Let $\mathcal{I}_2^\dagger = \{4, 7\}$. Then $\mathcal{L}_2 = \{4, 7, 12, 15\}$ is the set of all elements in \mathcal{I}_2 that leave a remainder of 4 mod 4 = 0 or 7 mod 4 = 3 on dividing by $2^{k_2+1} = 4$. Removing such numbers from \mathcal{I}_2 , we have $\mathcal{I}_3 = \mathcal{I}_2 \setminus \mathcal{L}_2 = \{6, 14\}$.
- 3) ($r = 3$) Now clearly $k_3 = 0$. Let $\mathcal{I}_3^\dagger = \{6\}$. Then $\mathcal{L}_4 = \{6, 14\}$, $\mathcal{I}_4 = \emptyset$ and the algorithm terminates.

According to the theorem, we have $|\Omega(\mathcal{I})| = 2^{k_1} + 2^{k_2} + 2^{k_3} = 7$, and an example $\Omega(\mathcal{I})$ is given by $\mathcal{I}_1^\dagger \cup \mathcal{I}_2^\dagger \cup \mathcal{I}_3^\dagger = \{0, 1, 2, 3, 4, 6, 7\}$.

We have several additional comments. First, we can say more about the formula for $|\Omega(\mathcal{I})|$. Since the k_r 's are nonincreasing, a typical sequence is, say,

$$\underbrace{l_1, l_1, \dots}_{\alpha_1 \text{ times}} \quad \underbrace{l_2, l_2, \dots}_{\alpha_2 \text{ times}} \quad \underbrace{l_3, l_3, \dots}_{\alpha_3 \text{ times}} \quad \dots$$

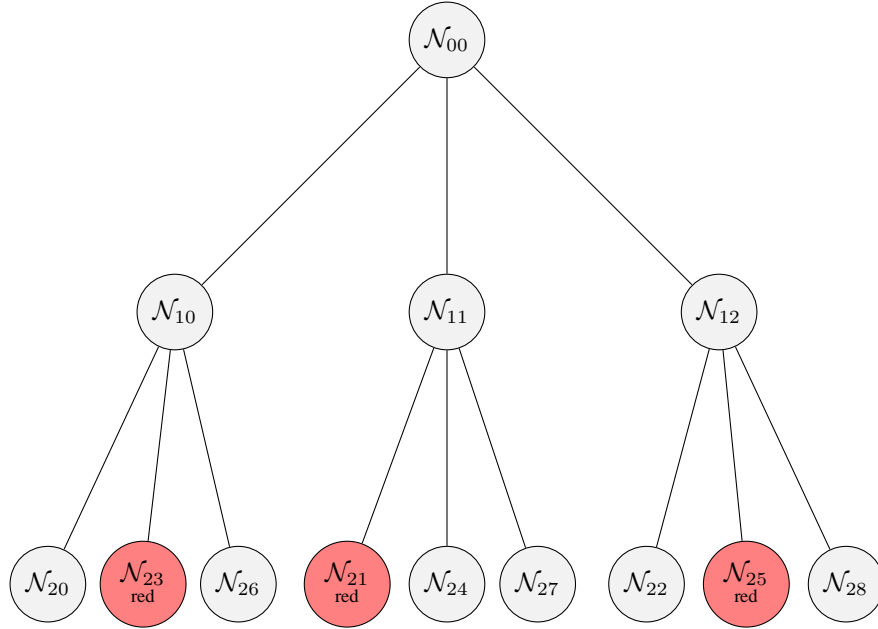


Fig. 4: Similar to Figure 3 but with root $\mathcal{N}_{00} = [0 : N - 1]$, this tree shows the relationship between \mathcal{N}_{ka} , for $p = 3$, $M \geq 3$.

with $l_1 > l_2 > l_3$. Given this, equation (30) appears as

$$|\Omega(\mathcal{I})| = \alpha_1 p^{l_1} + \alpha_2 p^{l_2} + \alpha_3 p^{l_3} + \dots \quad (32)$$

In fact, effectively, Theorem 6 constructs a base p expansion of $|\Omega(\mathcal{I})|$ because each power of p appears at most $p-1$ times.

Corollary 2: Let $\mathcal{I} \subseteq [0 : p^M - 1]$. The formula (32) is of the form,

$$|\Omega(\mathcal{I})| = \sum_r p^{k_r} = \sum_s \alpha_s p^{l_s},$$

with $l_1 > l_2 > l_3 > \dots$ and $\alpha_s \in [0 : p-1]$ for all s .

Proof: Begin with $\Omega(\mathcal{I}) = \bigcup_r \mathcal{I}_r^\dagger$. Since the \mathcal{I}_r^\dagger are disjoint, we have

$$\begin{aligned} \sum_{r=1}^{\alpha_1} \tilde{\chi}_{l_1+1}(a; \mathcal{I}_r^\dagger) &= \tilde{\chi}_{l_1+1}(a; \bigcup_{r=1}^{\alpha_1} \mathcal{I}_r^\dagger) \\ &\leq \tilde{\chi}_{l_1+1}(a; \Omega(\mathcal{I})), \quad a \in [0 : p^{l_1+1} - 1]. \end{aligned}$$

Summing this over all $a \in [0 : p^{l_1+1} - 1]$ we have

$$\begin{aligned} \alpha_1 p^{l_1} &= \sum_{r=1}^{\alpha_1} |\mathcal{I}_r^\dagger| = \sum_{r=1}^{\alpha_1} \sum_a \tilde{\chi}_{l_1+1}(a; \mathcal{I}_r^\dagger) \\ &\leq \sum_i \tilde{\chi}_{l_1+1}(a; \Omega(\mathcal{I})) = |\Omega(\mathcal{I})| < p^{l_1+1}, \end{aligned} \quad (33)$$

so $\alpha_1 < p$. For the last inequality in (33) we have used the upper bound from part (ii) in Theorem 5. We have also used that $|\mathcal{I}_r| = p^{k_r}$. The proof for other α_s is similar. For example, to prove that $\alpha_2 < p$ we start with

$$\sum_{r=\alpha_1+1}^{\alpha_2} \tilde{\chi}_{l_2+1}(a; \mathcal{I}_r^\dagger) \leq \tilde{\chi}_{l_2+1}(a; \Omega(\mathcal{I}_{\alpha_1+1}))$$

instead of (33). ■

If the algorithm above were initialized with a universal set \mathcal{I} , then from Theorem 6 we would obtain $\mathcal{I} = \Omega(\mathcal{I}) = \bigcup_r \mathcal{I}_r^\dagger$. This allows us to conclude that any universal set \mathcal{I} is a union

of elementary universal sets. Moreover, the sets \mathcal{I}_r^\dagger defined by the algorithm satisfy conditions in Lemma 7. For condition (ii), note that in the algorithm the set \mathcal{I}_r is recursively defined as $\mathcal{I}_r = \mathcal{I}_{r-1} \setminus \mathcal{L}_{r-1}$, so that

$$\mathcal{I}_r = (((\mathcal{I} \setminus \mathcal{L}_1) \setminus \mathcal{L}_2) \dots \setminus \mathcal{L}_{r-1}) = \mathcal{I} \setminus \left(\bigcup_{j=1}^{r-1} \mathcal{L}_j \right).$$

Hence $\mathcal{I}_r \cap \left(\bigcup_{j=1}^{r-1} \mathcal{L}_j \right) = \emptyset$. Then the sets \mathcal{I}_r^\dagger , obtained by the algorithm, satisfy $\mathcal{I}_r^\dagger \cap \left(\bigcup_{j=1}^{r-1} \mathcal{L}_j \right) = \emptyset$, since $\mathcal{I}_r^\dagger \subseteq \mathcal{I}_r$. Putting all these comments together we have the converse of Lemma 7, and then adding Theorem 6 we can state

Corollary 3: $\mathcal{I} \subseteq [0 : p^M - 1]$ is universal if and only if there exist

- (i) A nonincreasing finite sequence $k_1 \geq k_2 \geq \dots \geq 0$, with each value of k_r repeating at most $p-1$ times;
- (ii) Sets $\mathcal{I}_r^\dagger \subseteq \mathcal{I}$ with $\mathcal{I} = \bigcup_r \mathcal{I}_r^\dagger$;

such that

- (iii) \mathcal{I}_r^\dagger is a k_r -elementary universal set;
- (iv) $\mathcal{I}_r^\dagger \cap \left(\bigcup_{j=1}^{r-1} \mathcal{L}_j \right) = \emptyset$, where

$$\mathcal{L}_j = \{x \in [0 : N-1] : x \equiv i \pmod{p^{k_j} + 1} \text{ for some } i \in \mathcal{I}_j^\dagger\}.$$

Note that from (i), (ii) and (iii) we can also conclude that

$$|\mathcal{I}| = \sum_r |\mathcal{I}_r^\dagger| = \sum_r p^{k_r},$$

so the k_r are the powers of p appearing in the base- p expansion of $|\mathcal{I}|$, taken with repetitions. For example with $N = 9$, $|\mathcal{I}| = 7 = 2 \cdot 3^1 + 1 \cdot 3^0$, we expect the universal set $\mathcal{I} = \mathcal{I}_1^\dagger \cup \mathcal{I}_2^\dagger \cup \mathcal{I}_3^\dagger$ with \mathcal{I}_1^\dagger and \mathcal{I}_2^\dagger being 1-elementary, and \mathcal{I}_3^\dagger being 0-elementary. Corollary 3 implies that the k_r read off from the base- p expansion of $|\mathcal{I}|$ must be the same as the k_r generated by the algorithm if \mathcal{I} is universal.

Remark 4 (Universal sets of prescribed order):

As it stands, the algorithm finds a universal set of the largest size contained in \mathcal{I} . With Corollary 3 we can now modify the algorithm to solve the following problem:

Given a set $\mathcal{I} \subseteq [0 : p^M - 1]$, and $d \leq |\Omega(\mathcal{I})|$, find a universal set $\mathcal{J} \subseteq \mathcal{I}$ with $|\mathcal{J}| = d$.

We follow the algorithm as in steps 1-4, but we change the definition of k_r in Step 1. Write the base- p expansion of d with repetitions, $d = \sum_r p^{k_r}$, read off the k_r as the powers of p that appear in the expansion, and arrange the k_r in nonincreasing order. This ensures that condition (i) in Corollary 3 is satisfied. The construction of the \mathcal{I}_r^\dagger in Steps 2-4 of the algorithm will ensure that (iii) and (iv) are satisfied. We conclude that with the \mathcal{I}_r^\dagger so obtained by the algorithm the set $\mathcal{J} = \bigcup_r \mathcal{I}_r^\dagger \subset \mathcal{I}$ is universal, and it is of the right size by definition of the k_r .

Finally, we have

Proof of Theorem 6: As observed above, the \mathcal{I}_r^\dagger generated by the algorithm satisfy the hypotheses of Lemma 7, so the set $\bigcup_r \mathcal{I}_r^\dagger$ is universal. If we show

$$|\Omega(\mathcal{I})| \leq \sum_r p^{k_r},$$

then Theorem 6 follows.

For this we prove

$$|\Omega(\mathcal{I}_r)| \leq p^{k_r} + |\Omega(\mathcal{I}_{r+1})|. \quad (34)$$

We appeal to Theorem 5 to find a maximal universal sampling set \mathcal{A} with $\mathcal{I}_r^\dagger \subseteq \mathcal{A} \subseteq \mathcal{I}_r$, and we will show

$$\mathcal{A} \setminus \mathcal{I}_r^\dagger \text{ is universal,} \quad (35)$$

$$\mathcal{A} \setminus \mathcal{I}_r^\dagger \subseteq \mathcal{I}_{r+1}. \quad (36)$$

Since $|\mathcal{A}| = |\Omega(\mathcal{I}_r)|$ These imply

$$|\Omega(\mathcal{I}_r)| - p^{k_r} = |\mathcal{A} \setminus \mathcal{I}_r^\dagger| \leq |\Omega(\mathcal{I}_{r+1})|,$$

which is (34).

First (35). Now,

$$\tilde{\chi}_s(a; \mathcal{A} \setminus \mathcal{I}_r^\dagger) = \tilde{\chi}_s(a; \mathcal{A}) - \tilde{\chi}_s(a; \mathcal{I}_r^\dagger), \quad a \in [0 : p^s - 1],$$

and for $s \leq k_r$ the second term is constant,

$$\tilde{\chi}_s(a; \mathcal{I}_r^\dagger) = p^{k_r - s},$$

from (24). Since \mathcal{A} is universal, $|\tilde{\chi}_s(a; \mathcal{A}) - \tilde{\chi}_s(b; \mathcal{A})| \leq 1$ for all $a, b \in [0 : p^s - 1]$ and for all s , so we at least have

$$|\tilde{\chi}_s(a; \mathcal{A} \setminus \mathcal{I}_r^\dagger) - \tilde{\chi}_s(b; \mathcal{A} \setminus \mathcal{I}_r^\dagger)| \leq 1,$$

for all $s \leq k_r$. We need to check that this inequality continues to hold for $s \geq k_r + 1$.

As we have argued before, by the definition of k_r at least one congruence class in \mathcal{I}_r/p^s is empty when $s \geq k_r + 1$, so $\tilde{\chi}_s(a_0; \mathcal{I}_r) = 0$ for some a_0 , and because $\mathcal{A} \subseteq \mathcal{I}_r$ is universal we have $\tilde{\chi}_s(a; \mathcal{A}) \leq 1$ for all a . Furthermore, $\mathcal{I}_r^\dagger \subseteq \mathcal{A}$ implies

$$0 \leq \tilde{\chi}_s(a; \mathcal{A}) - \tilde{\chi}_s(a; \mathcal{I}_r^\dagger) = \tilde{\chi}_s(a; \mathcal{A} \setminus \mathcal{I}_r^\dagger).$$

Hence the values of $\tilde{\chi}_s(a; \mathcal{A} \setminus \mathcal{I}_r^\dagger)$ are in $\{0, 1\}$ and consequently

$$|\tilde{\chi}_s(a; \mathcal{A} \setminus \mathcal{I}_r^\dagger) - \tilde{\chi}_s(b; \mathcal{A} \setminus \mathcal{I}_r^\dagger)| \leq 1,$$

for all $s \geq k_r + 1$. This establishes that $\mathcal{A} \setminus \mathcal{I}_r^\dagger$ is universal.

We prove (36) by contradiction. If it were not true that $\mathcal{A} \setminus \mathcal{I}_r^\dagger \subseteq \mathcal{I}_{r+1}$ then there would exist an $x \in (\mathcal{A} \setminus \mathcal{I}_r^\dagger) \cap \mathcal{I}_r$. Then $\tilde{\chi}_{k_r+1}([x]_{k_r+1}; \mathcal{A}) = 2$, for on dividing by p^{k_r+1} , x leaves a remainder of $[x]_{k_r+1}$ (by definition) and so does one other element in \mathcal{I}_r^\dagger . But this contradicts $\tilde{\chi}_s(a; \mathcal{A}) \leq 1$ for $s \geq k_r + 1$ from the preceding paragraph.

This completes the proof of Theorem 6. \blacksquare

Remark 5: We can give an upper bound for the computational complexity of the algorithm for constructing a universal sampling set of size d (including constructing a maximal universal sampling set). Within an iteration, in the worst case the algorithm makes a complete pass over all the nodes of the tree once, and the number of nodes is $O(N)$. Further, the number of iterations is $\alpha_1 + \alpha_2 + \dots + \alpha_M$ where

$$d = \alpha_1 p^{M-1} + \alpha_2 p^{M-2} + \dots + \alpha_{M-1} p + \alpha_M.$$

Hence the largest number of iterations is $(p-1)M$, and the complexity of the algorithm is at most $O(N \log N)$.

D. Counting Universal Sets

The preceding structure theorems allow us to find the number of universal sampling sets $\mathcal{I} \subseteq [0 : p^M - 1]$ of size d . The formula uses the digits from the base- p expansion of d , and as above we let

$$d = \alpha_1 p^{M-1} + \alpha_2 p^{M-2} + \dots + \alpha_{M-1} p + \alpha_M,$$

where $0 \leq \alpha_i < p$. For $i = 0, 1, \dots, M$ define

$$d_i = \sum_{j=i+1}^M \alpha_j p^{M-j}.$$

Hence $d_0 = d$ and $d_M = 0$.

Theorem 7: The number of universal sampling sets in $[0 : p^M - 1]$ of size d is

$$\mathcal{C}(d, p^M) = \prod_{i=1}^M \binom{p}{\alpha_i + 1}^{d_i} \binom{p}{\alpha_i}^{p^{M-i} - d_i}.$$

Proof: The proof goes by establishing a recurrence relation for \mathcal{C} in the d_i .² Let \mathcal{I} be a universal sampling set of size d and construct the congruence tree as in Figure 2 with root $\mathcal{I}_{00} = \mathcal{I}$. We first note that d_1 of the nodes at level $M-1$ have weight $\alpha_1 + 1$ and the remaining $p^{M-1} - d_1$ nodes have weight α_1 , where

$$d_1 = \sum_{i=2}^M \alpha_i p^{M-i}.$$

The proof for this is along the same lines as the argument in the proof of Theorem 4, (i) \iff (ii). Figure 5 illustrates this. The singleton blue nodes at the bottom level are the elements of \mathcal{I} , and the other nodes (which would be the singletons $\{6\}$ and $\{7\}$) are empty. The red nodes at the penultimate level represent the nodes that have weight $\alpha_1 + 1$ (and there are d_1 of them).

²We are grateful to a reviewer for suggesting a way to make greater use of the recursive aspect of our original argument, resulting in a much shorter and cleaner proof.

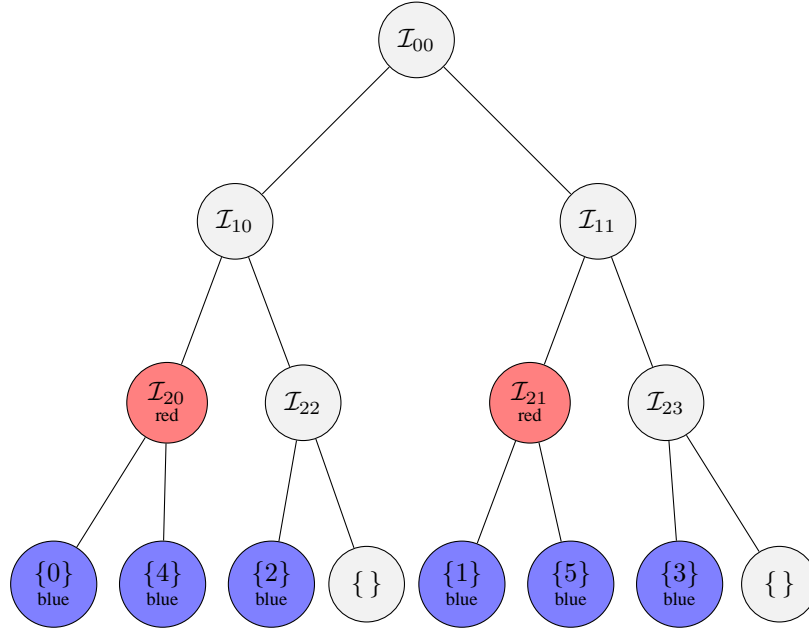


Fig. 5: The congruence class tree for $N = 8$. The universal sampling set $\{0, 1, 2, 3, 4, 5\}$ of size $d = 6$ is represented by the blue nodes at the bottom level. The red nodes at the penultimate level represent the nodes that have weight 2, the rest of the nodes at the penultimate level have weight 1.

Now remove the bottom level of the tree, effectively making $N = p^{M-1}$, and resulting in Figure 6. If the starting set (the blue nodes in Figure 5) is universal, then so must be the set formed by the red nodes in Figure 6. Hence the number of ways of choosing the red nodes is the same as the number of universal sampling sets of size d_1 in $[0 : p^{M-1} - 1]$, that is $\mathcal{C}(d_1, p^{M-1})$.

Once the red nodes are chosen, we need to choose the blue nodes by taking $\alpha_1 + 1$ elements from the red nodes and α_1 elements from the remaining (non-red) nodes, which can be done in

$$\binom{p}{\alpha_1 + 1}^{d_1} \binom{p}{\alpha_1}^{p^{M-1} - d_1}$$

ways. Hence

$$\mathcal{C}(d, p^M) = \binom{p}{\alpha_1 + 1}^{d_1} \binom{p}{\alpha_1}^{p^{M-1} - d_1} \mathcal{C}(d_1, p^{M-1}). \quad (37)$$

This full formula follows. ■

One special case of the counting formula is easy to evaluate.

Corollary 4: Let $d = p^k$ where $k < M$. Then the number of universal sets of size d in $[0 : p^M - 1]$ is $(p^M/d)^d$.

In particular when $N = 2^M$, and $d = 2^{M-1} = N/2$, the number of universal sets is $2^{N/2}$. On the other hand, the total number of sets of size 2^{M-1} in $[0 : 2^N - 1]$ is

$$\binom{N}{N/2} \approx 2^N / \sqrt{\pi N}$$

by Stirling's approximation. Hence the fraction of sets that are universal is approximately $\sqrt{\pi N} / 2^{N/2}$, which decreases exponentially with N .

The function $\mathcal{C}(d, p^M)$ is certainly complicated, but it has some remarkable properties. Though not clear from the




Fig. 6: Remove the bottom level of the tree in Figure 5. The resulting red nodes are a universal sampling set in $[0 : 3]$

formula, we have

$$\mathcal{C}(d, p^M) = \mathcal{C}(p^M - d, p^M).$$

This follows from the following lemma, which is itself a simple but interesting property of universal sampling sets.

Lemma 8: If $\mathcal{A} \subseteq [0 : p^M - 1]$ is a universal sampling set then so is $\mathcal{A}' = [0 : p^M - 1] \setminus \mathcal{A}$.

This extends the bracelet property of universal sampling sets, though for bracelets we need not assume that N is a prime power.

Proof: For any $0 \leq k \leq M$ and $a \in [0 : p^k - 1]$,

$$\begin{aligned} \tilde{\chi}_k(a; \mathcal{A}') &= \tilde{\chi}_k(a; [0 : p^M - 1]) - \tilde{\chi}_k(a; \mathcal{A}) \\ &= p^{M-k} - \tilde{\chi}_k(a; \mathcal{A}). \end{aligned}$$

Next, since

$$|\tilde{\chi}_k(a; \mathcal{A}) - \tilde{\chi}_k(b; \mathcal{A})| \leq 1,$$

for all $a, b \in [0 : p^k - 1]$, it follows that

$$|\tilde{\chi}_k(a; \mathcal{A}') - \tilde{\chi}_k(b; \mathcal{A}')| \leq 1.$$

Figure 7 displays $\log \mathcal{C}(d, 5^M)$ as a function of d as M takes increasing values. The plots show the symmetry, $\mathcal{C}(d, p^M) = \mathcal{C}(p^M - d, p^M)$, but they show much more. We can observe the following:

- (i) There are a series of bumps on several (visible) scales. One cannot fail to notice that at each scale the number of bumps in the graph is 5, which is the prime p here. Experiments with other primes have similar plots and in each case indicate that the number of bumps is equal to the prime.
- (ii) With increasing M the plots of the count are somehow converging in shape – they all start to look similar.

The second point can indeed be quantified. One can show that for each $\alpha \in [0, 1]$,

$$\lim_{M \rightarrow \infty} \frac{\log \mathcal{C}(\lfloor \alpha p^M \rfloor, p^M)}{p^M}$$

exists. See [7]. This compares nicely with the fact that a similar function with $\mathcal{C}(d, N)$ replaced by $\binom{N}{d}$ also converges, and to the entropy function:

$$\lim_{M \rightarrow \infty} \left(\frac{1}{p^M} \log \left(\frac{p^M}{\lfloor \alpha p^M \rfloor} \right) \right) = \alpha \log \frac{1}{\alpha} + (1 - \alpha) \log \frac{1}{1 - \alpha} = H(\alpha).$$

This is the limiting case of counting *all* index sets.

Plots of

$$\mathcal{H}_p(\alpha) = \lim_{M \rightarrow \infty} \frac{\log \mathcal{C}(\lfloor \alpha p^M \rfloor, p^M)}{p^M}, \quad 0 \leq \alpha \leq 1, \quad (38)$$

are shown in Figures 8 and 9 for several values of p , along with a plot of $H(\alpha)$.

The plots of $\mathcal{H}_p(\alpha)$ seem to satisfy observation (i), that the curves have p bumps at each scale. Here is an explanation. In the notation of Theorem 7, suppose $\alpha_1 = 0$ (i.e., $d < p^{M-1}$). Then $d_1 = d$ and we have, as in (37),

$$\begin{aligned} \mathcal{C}(d, p^M) &= \binom{p}{1}^{d_1} \binom{p}{0}^{p^{M-1}-d_1} \mathcal{C}(d_1, p^{M-1}) \\ &= p^d \mathcal{C}(d, p^{M-1}). \end{aligned}$$

Let $M \rightarrow \infty$, so $d/p^M \rightarrow \alpha$ (with $\alpha < p$). Then with reference to (38),

$$\mathcal{H}_p(\alpha) = \lim_{M \rightarrow \infty} \left(\frac{d}{p^M} \log p + \mathcal{H}_p(p\alpha) \right) = \alpha \log p + \mathcal{H}_p(p\alpha),$$

leading to the self-similar plots we observe.

E. Maximal and Minimal Universal Sampling Sets

Along with maximal universal sets is the allied notion of minimal universal sets.

Definition 7: Let $\mathcal{I} \subseteq [0 : N - 1]$. A *minimal universal sampling set* for \mathcal{I} is a universal sampling set of smallest cardinality that contains \mathcal{I} .

Again we need a notation and we let $\Phi(\mathcal{I})$ denote a generic minimal universal sampling set containing \mathcal{I} . Thus $|\Phi(\mathcal{I})| \leq |\mathcal{J}|$ for any universal sampling set $\mathcal{J} \supseteq \mathcal{I}$.

Let us show one way that maximal and minimal universal sampling sets are related. The proof relies on Lemma 8 from the previous subsection.

Theorem 8: Let $\mathcal{I} \subset [0 : p^M - 1]$, $\mathcal{I}' = [0 : p^M - 1] \setminus \mathcal{I}$. Then

$$|\Phi(\mathcal{I})| = p^M - |\Omega(\mathcal{I}')|.$$

Proof: Let $\mathcal{A}' = [0 : p^M - 1] \setminus \Phi(\mathcal{I})$. Then \mathcal{A}' is universal by Lemma 8. Since $\Phi(\mathcal{I}) \supseteq \mathcal{I}$ we have $\mathcal{A}' \subset [0 : N - 1] \setminus \mathcal{I} = \mathcal{I}'$ and hence

$$p^M - |\Phi(\mathcal{I})| = |\mathcal{A}'| \leq |\Omega(\mathcal{I}')|.$$

Similarly, let $\mathcal{B}' = [0 : p^M - 1] \setminus \Omega(\mathcal{I}')$. Then \mathcal{B}' is universal, it contains $[0 : p^M - 1] \setminus \mathcal{I}' = \mathcal{I}$ and so

$$p^M - |\Omega(\mathcal{I}')| = |\mathcal{B}'| \geq |\Phi(\mathcal{I})|.$$

Taken together the two inequalities prove the theorem. ■

VI. AN UNCERTAINTY PRINCIPLE, RANDOM SIGNALS, AND SUMSETS

Generally speaking, an “uncertainty principle” is an inequality relating the supports of a nonzero function and its Fourier transform, in the present setting $f: \mathbb{Z}_N \rightarrow \mathbb{C}$, and $\mathcal{F}f: \mathbb{Z}_N \rightarrow \mathbb{C}$. The notions of maximal and minimal universal sampling sets lead immediately to an additive uncertainty principle. Without the language of universality, Tao [6] made this connection in the case when N is a prime using Chebotarev’s theorem, see Corollary 5, though, as he states, it was probably already known as a folk theorem.

Let

$$\mathcal{Z}(f) = \{i: f(i) = 0\}$$

be the zero set of f . The support is the complement of the zero set, and we denote it by $\text{supp}(f)$. Our result is

Theorem 9: If f is not the zero function then

$$\begin{aligned} |\text{supp}(\mathcal{F}f)| &\geq 1 + |\Omega(\mathcal{Z}(f))|, \\ |\text{supp}(f)| &\geq 1 + |\Omega(\mathcal{Z}(\mathcal{F}f))|; \end{aligned} \quad (39)$$

and

$$\begin{aligned} |\mathcal{Z}(\mathcal{F}f)| + 1 &\leq |\Phi(\text{supp}(f))|, \\ |\mathcal{Z}(f)| + 1 &\leq |\Phi(\text{supp}(\mathcal{F}f))|. \end{aligned} \quad (40)$$

We are not assuming that N is a prime power here. However, we immediately deduce

Corollary 5 (Tao): If N is prime and f is not the zero function then

$$|\text{supp}(\mathcal{F}f)| + |\text{supp}(f)| \geq N + 1.$$




Fig. 7: Plots of $\log \mathcal{C}(d, p^M)$ vs d for powers of $p = 5$. Note the 5 bumps on different scales.




Fig. 8: Plots of the limit of the counting functions for $p = 2, 5$ compared to the Entropy function. Note the self-similarity as it depends on the prime.




Fig. 9: Similar to Figure 8 with $p = 3, 7$

Proof: If N is prime, then by Chebotarev's theorem every index set is universal. In particular the set $\mathcal{Z}(f)$ is universal. Hence $\Omega(\mathcal{Z}(f)) = \mathcal{Z}(f)$. From Theorem 9,

$$\begin{aligned} |\text{supp}(\mathcal{F}f)| &\geq 1 + |\Omega(\mathcal{Z}(f))| \\ &= 1 + |\mathcal{Z}(f)| = 1 + N - |\text{supp}(f)|. \end{aligned}$$

We also have

Corollary 6: Suppose f vanishes on a set of consecutive integers \mathcal{I} . Then $|\text{supp}(\mathcal{F}f)| \geq |\mathcal{I}| + 1$. If \mathcal{J} is a set of integers such that $\mathcal{F}f(\mathcal{J}) = 0$, then $|\mathcal{I}| + |\mathcal{J}| \leq N - 1$.

Proof: We observed previously that any set of consecutive integers, \mathcal{I} in this case, is universal. Since $\mathcal{I} \subseteq \mathcal{Z}(f)$, we have $|\Omega(\mathcal{Z}(f))| \geq |\mathcal{I}|$. From Theorem 9, this implies $|\text{supp}(\mathcal{F}f)| \geq |\mathcal{I}| + 1$. Further, if $\mathcal{F}f(\mathcal{J}) = 0$ then $N - |\mathcal{J}| \geq |\text{supp}(\mathcal{F}f)|$ and so $N - |\mathcal{J}| \geq |\mathcal{I}| + 1$. ■

The proof of Theorem 9 itself is very brief.

Proof of Theorem 9: Suppose $|\text{supp}(\mathcal{F}f)| \leq |\Omega(\mathcal{Z}(f))|$. From $\Omega(\mathcal{Z}(f)) \subseteq \mathcal{Z}(f)$ it follows that f vanishes on $\Omega(\mathcal{Z}(f))$. Since $\Omega(\mathcal{Z})$ is a universal sampling set this implies that $\mathcal{F}f \equiv 0$, contradicting the assumption that f is not the zero function. This proves the first statement in (39). A similar argument establishes the second statement.

For the proof of (40), write $\mathcal{Z} = \mathcal{Z}(\mathcal{F}f)$ and $\mathcal{A} = \Phi(\text{supp}(f))$. Then

$$\mathcal{F}f(\mathcal{Z}) = 0 \text{ and so } E_{\mathcal{Z}}^T \mathcal{F}f = 0.$$

However f is supported within \mathcal{A} , and so we may write $f = E_{\mathcal{A}}g$, where $g = f(\mathcal{A}) \neq 0$. This means we must have

$$E_{\mathcal{Z}}^T \mathcal{F} E_{\mathcal{A}} g = 0, \text{ for some } g \neq 0, \quad (41)$$

i.e. the columns of $E_{\mathcal{Z}}^T \mathcal{F} E_{\mathcal{A}}$ are dependent. This is expected if $|\mathcal{Z}| < |\mathcal{A}|$. However, if $|\mathcal{Z}| \geq |\mathcal{A}|$, this contradicts the universality of \mathcal{A} . Hence we must have $|\mathcal{Z}| \leq |\mathcal{A}| - 1$, which is the first inequality in (40). A similar argument establishes the second statement. ■

It is interesting that when N is a prime power the two statements (39) and (40) are equivalent. To see this we first derive (40) from (39) when $N = p^M$. This appeals to Theorem 8 on the relation between maximal and minimal sets, with $\text{supp}(f) = [0 : N - 1] \setminus \mathcal{Z}(f)$. Thus, from (39), $|\text{supp}(\mathcal{F}f)| \geq 1 + |\Omega(\mathcal{Z}(f))|$, and substituting from Theorem 8,

$$|\text{supp}(\mathcal{F}f)| \geq 1 + N - |\Phi(\text{supp}(f))|.$$

But $|\text{supp}(\mathcal{F}f)| = N - |\mathcal{Z}(\mathcal{F}f)|$, so

$$N - |\mathcal{Z}(\mathcal{F}f)| \geq 1 + N - |\Phi(\text{supp}(f))|,$$

which is the same as the first statement in (40). Again, the second statement in (40) follows in a similar manner. We could have started instead with (40) and from this derived (39).

In cases where $\mathcal{Z}(f)$ itself is a universal sampling set, the uncertainty principle in Theorem 9 can be as strong as the uncertainty principle for the prime N case.

Remark 6: Readers familiar with the seminal paper of Donoho and Stark [14] will wonder if the additive uncertainty principle in Theorem 9 can be applied to the problem of reconstruction of a signal corrupted by sparse noise. (See also [15] for more recent work.) The answer is yes, and we refer to [16].

A. Random Index Sets and Random Signals

We will give several applications of these ideas. First we combine Theorem 9 with a probabilistic estimate on the size of a maximal universal sampling set for randomly chosen index sets. We must revert to the assumption that N is a prime power.

Theorem 10: Let $N = p^M$. Let \mathcal{R}_s be an index set of s numbers chosen at random from $[0 : N - 1]$. Let $\lambda = (N - s)/N$. If $d, \delta > 0$ satisfy

$$N \log(1/\lambda) \geq (1 + \delta)d \log d, \quad (42)$$

then $|\Omega(\mathcal{R}_s)| \geq d$ with probability at least $1 - d^{-\delta}$.

This means that if we can choose a large d satisfying (42), which is possible, for example, if N is large and λ is small, then $|\Omega(\mathcal{R}_s)| \geq d$ with high probability. Thus while it is unlikely that a randomly chosen index set will be universal, it is quite likely that such an index set will contain a large universal set as a subset.

We will apply Theorem 10 to the case when \mathcal{R}_s is the zero set of $f: \mathbb{Z}_N \rightarrow \mathbb{C}$. Then $\lambda = |\text{supp}(f)|/N$, i.e., λ is the fraction of nonzero entries in f .

Proof: The proof uses the bound in part (ii) of Theorem 5. Let k be the largest integer such that no congruence classes in \mathcal{R}_s/p^k are empty. Note that k is random since \mathcal{R}_s is random. Then $|\Omega(\mathcal{R}_s)| \leq d-1$ implies

$$p^k \leq |\Omega(\mathcal{R}_s)| \leq d-1,$$

by Theorem 5. Therefore

$$\begin{aligned} \text{Prob}(|\Omega(\mathcal{R}_s)| \leq d-1) &\leq \text{Prob}(p^k \leq d-1) \\ &= \text{Prob}(k \leq \lfloor \log_p(d-1) \rfloor) \\ &= \text{Prob}(\text{at least one congruence class in} \\ &\quad \mathcal{R}_s/p^{\lfloor \log_p(d-1) \rfloor + 1} \text{ is empty}). \end{aligned} \quad (43)$$

We will compute the last probability.

Let $b = \lfloor \log_p(d-1) \rfloor + 1$, and let \mathcal{N}_{ba} be the set of elements in $[0 : N-1]$ that leave a remainder of $a \in [0 : p^b-1]$ when divided by p^b . Since $N = p^M$ all of the \mathcal{N}_{ba} have size $t = N/p^b = p^{M-b}/p^{\lfloor \log_p(d-1) \rfloor + 1}$.

Fix a particular residue a . The probability that $\mathcal{N}_{ba} \cap \mathcal{R}_s$ is empty (in words, the probability that a particular congruence class goes missing in \mathcal{R}_s) is $\binom{N-t}{s} / \binom{N}{s}$. This is because the number of ways of picking \mathcal{R}_s is $\binom{N}{s}$ while the number of ways of picking \mathcal{R}_s so that $\mathcal{N}_{ba} \cap \mathcal{R}_s = \emptyset$ is the number of ways of picking s elements from

$$|[0 : N-1] \setminus \mathcal{N}_{ba}| = N-t$$

elements. Then

$$\begin{aligned} \text{Prob}(\mathcal{N}_{ba} \cap \mathcal{R}_s = \emptyset) &= \frac{\binom{N-t}{s}}{\binom{N}{s}} \\ &= \frac{(N-t-s)(N-t-2-s)\dots(N-s)}{(N-t+1)(N-t+2)\dots N} \\ &= \left(1 - \frac{s}{N-t+1}\right) \left(1 - \frac{s}{N-t+2}\right) \dots \left(1 - \frac{s}{N}\right) \\ &\leq \left(1 - \frac{s}{N}\right) \left(1 - \frac{s}{N}\right) \dots \left(1 - \frac{s}{N}\right) = \left(1 - \frac{s}{N}\right)^t. \end{aligned} \quad (44)$$

From this,

$$\begin{aligned} \text{Prob}(\text{at least one congruence class in} \\ \mathcal{R}_s/p^{\lfloor \log_p(d-1) \rfloor + 1} \text{ is empty}) \\ &= \text{Prob}\left(\bigcup_i (\mathcal{N}_{ba} \cap \mathcal{R}_s = \emptyset)\right) \\ &\leq \sum_i \text{Prob}(\mathcal{N}_{ba} \cap \mathcal{R}_s = \emptyset) \\ &\leq \frac{N}{t} \left(1 - \frac{s}{N}\right)^t = N\lambda^t/t. \end{aligned} \quad (45)$$

Hence we have from (45),

$$\text{Prob}(|\Omega(\mathcal{R}_s)| \leq d-1) \leq N\lambda^t/t. \quad (46)$$

Now, $t = N/p^{\lfloor \log_p(d-1) \rfloor + 1} \geq N/d$, since $\lfloor x \rfloor \leq x$. Using this in (46),

$$\begin{aligned} \text{Prob}(|\Omega(\mathcal{R}_s)| \leq d-1) &\leq N\lambda^t/t \leq d\lambda^{N/d} \\ &= \exp\left(\log d - \frac{N \log(1/\lambda)}{d}\right) \\ &= \exp\left(\log d \left(1 - \frac{N \log(1/\lambda)}{d \log d}\right)\right) \\ &\leq \exp(-\delta \log d) \quad (\text{from the hypothesis of the theorem}) \\ &= d^{-\delta}. \end{aligned} \quad (47)$$

We conclude that $\text{Prob}(|\Omega(\mathcal{R}_s)| \geq d) \geq 1 - d^{-\delta}$. ■

We can now state a probabilistic uncertainty principle. Afterward we will comment on how this compares to the result of Candes, Romberg and Tao [3].

Theorem 11: Let $N = p^M$. Let $\mathcal{G}_{N,r}$ be the set of all signals $g: \mathbb{Z}_N \rightarrow \mathbb{C}$ with support of size r . Let $g \in \mathcal{G}_{N,r}$ be a signal whose support is drawn at random from the set of all index sets of size r . Let the values of g on the support set be drawn according to some arbitrary distribution. For $\delta > 0$ let

$$a_{N,\delta} = \frac{N}{(1+\delta) \log N} (1 + \log(1+\delta) + \log \log N).$$

Then

$$|\text{supp}(g)| + |\text{supp}(\mathcal{F}g)| \geq 1 + a_{N,\delta} \quad (48)$$

with probability at least $1 - (a_{N,\delta} - r)^{-\delta}$.

If r is small compared to $a_{N,\delta}$, Theorem 11 states that almost all signals g in $\mathcal{G}_{N,r}$ satisfy the uncertainty principle above; roughly speaking

$$|\text{supp}(g)| + |\text{supp}(\mathcal{F}g)| \geq N(1 + \log \log N) / \log N$$

for most g .

Proof: Picking the support of g at random among sets of size r is equivalent to picking the zero set of g at random among all index sets of size $N-r$. The proof now makes use of Theorem 10 to get a lower bound on $|\Omega(\mathcal{Z}(g))|$. For this we need to choose d, δ so that

$$N \log(1/\lambda) = N \log N/r > (1+\delta)d \log d. \quad (49)$$

Fix any $\delta > 0$ and let $d = N \log(N/r)/(1+\delta) \log N$. We check that d, δ satisfy (49):

$$\begin{aligned} (1+\delta)d \log d &= \frac{N \log(N/r)}{\log N} \log \left(\frac{N \log(N/r)}{(1+\delta) \log N} \right) \\ &< \frac{N \log(N/r)}{\log N} \log N = N \log N/r, \end{aligned}$$

Then from Theorem 10,

$$|\Omega(\mathcal{Z}(g))| \geq N \log(N/r)/(1+\delta) \log N$$

with probability $1 - d^{-\delta}$. From the uncertainty principle Theorem 9, we now have

$$\begin{aligned} |\text{supp}(\mathcal{F}g)| &\geq 1 + |\Omega(\mathcal{Z}(g))| \\ &\geq 1 + N \log(N/r)/(1+\delta) \log N \end{aligned}$$

with probability $1 - d^{-\delta}$.

The final step in the proof uses a lower bound on $d = N \log(N/r)/(1 + \delta) \log N$. We have set apart this technical result as Lemma 9, below. This gives

$$|\text{supp}(\mathcal{F}g)| \geq 1 + a_{N,\delta} - r$$

with probability $1 - d^{-\delta}$. Since $1 - d^{-\delta} \geq 1 - (a_{N,\delta} - r)^{-\delta}$, we can say

$$|\text{supp}(\mathcal{F}g)| \geq 1 + a_{N,\delta} - r$$

with probability $1 - (a_{N,\delta} - r)^{-\delta}$. The result follows since $r = |\text{supp}(g)|$. ■

Lemma 9: Let

$$d = \frac{N \log(N/r)}{(1 + \delta) \log N}$$

and

$$a_{N,\delta} = \frac{N}{(1 + \delta) \log N} (1 + \log(1 + \delta) + \log \log N),$$

as in Theorem 11. Then $d \geq a_{N,\delta} - r$.

Proof: The convex function $\log(N/r)$ is bounded below by its tangent at any point $r_0 > 0$. Thus

$$\log(N/r) \geq \log(N/r_0) + \left(-\frac{1}{r_0}(r - r_0)\right).$$

For

$$r_0 = \frac{N}{(1 + \delta) \log N},$$

this reads

$$\begin{aligned} \log(N/r) &\geq \log((1 + \delta) \log N) \\ &\quad + \left(-\frac{(1 + \delta) \log N}{N} \left(r - \frac{N}{(1 + \delta) \log N}\right)\right). \end{aligned}$$

Multiplying by $N/(1 + \delta) \log N$, we have

$$\begin{aligned} d &= \frac{N \log(N/r)}{(1 + \delta) \log N} \\ &\geq \frac{N \log((1 + \delta) \log N)}{(1 + \delta) \log N} - \left(r - \frac{N}{(1 + \delta) \log N}\right) \\ &= \frac{N}{(1 + \delta) \log N} (\log(1 + \delta) + 1 + \log \log N) - r \\ &= a_{N,\delta} - r. \end{aligned}$$

Remark 7: The robust uncertainty principle of Candes, Romberg and Tao in [3] is as follows: for $M > 0$ there exists a constant C_M such that

$$|\text{supp}(g)| + |\text{supp}(\mathcal{F}g)| \geq C_M N (\log N)^{-1/2},$$

with probability $1 - O(N^{-M})$. This inequality is stronger than that of Theorem 11 by about $(\log N)^{-1/2}$. Also, Theorem 11 holds for $N = p^M$, whereas the inequality above holds for all N .

In our proof of Theorem 10 we have only used the bound $|\Omega(\mathcal{Z}(g))| \geq p^k$ from Theorem 5. By using the exact formula for $|\Omega(\mathcal{Z}(g))|$ in Theorem 6 (or by a better lower bound) it might be possible to tighten the uncertainty principle of Theorem 11 and remove the factor $(\log N)^{-1/2}$.

B. Sumsets and the Cauchy-Davenport Theorem

Our final application is a generalization of the Cauchy-Davenport theorem [17], from additive number theory, on the size of sumsets. Again the inspiration comes from Tao's approach, [6], to the original Cauchy-Davenport theorem via Chebotarev's theorem.

Theorem 12: Let $\mathcal{X}, \mathcal{Y} \subseteq [0 : N - 1]$. If either \mathcal{X} or \mathcal{Y} is a universal sampling set, then

$$|\mathcal{X} + \mathcal{Y}| \geq |\mathcal{X}| + |\mathcal{Y}| - 1, \quad (50)$$

when $|\mathcal{X}| + |\mathcal{Y}| - 1 \leq N$.

Here $\mathcal{X} + \mathcal{Y}$ is the sumset defined as

$$\mathcal{X} + \mathcal{Y} = \{x + y : x \in \mathcal{X}, y \in \mathcal{Y}\},$$

where the addition is modulo N .

We are not assuming that N is a prime power, while the classical theorem has $N = p$ and there are no assumptions on \mathcal{X} or \mathcal{Y} . That form of the result follows from Theorem 12, since all index sets in $[0 : N - 1]$ are universal when N is prime.

As a corollary we get a statement on the size of $|\mathcal{X} + \mathcal{Y}|$ without making an assumption on \mathcal{X} or \mathcal{Y} .

Corollary 7: Let $\mathcal{X}, \mathcal{Y} \subseteq [0 : N - 1]$ be index sets. Then,

$$|\mathcal{X} + \mathcal{Y}| \geq \max\{|\Omega(\mathcal{X})| + |\mathcal{Y}| - 1, |\mathcal{X}| + |\Omega(\mathcal{Y})| - 1\}. \quad (51)$$

Proof: Since $\Omega(\mathcal{X}) \subseteq \mathcal{X}$, it follows that $\Omega(\mathcal{X}) + \mathcal{Y} \subseteq \mathcal{X} + \mathcal{Y}$. Now,

$$|\mathcal{X} + \mathcal{Y}| \geq |\Omega(\mathcal{X}) + \mathcal{Y}| \geq |\Omega(\mathcal{X})| + |\mathcal{Y}| - 1 \text{ from Theorem 12.}$$

The inequality $|\mathcal{X} + \mathcal{Y}| \geq |\mathcal{X}| + |\Omega(\mathcal{Y})| - 1$ follows similarly. ■

Proof of Theorem 12: First note that (50) follows trivially when either \mathcal{X} or \mathcal{Y} is a singleton. (More precisely, if, say, \mathcal{X} is a singleton, then $\mathcal{X} + \mathcal{Y}$ is just a translate of \mathcal{Y} , and so (50) holds with equality). For the rest of the proof, we assume that $|\mathcal{X}|, |\mathcal{Y}| \geq 2$. Let $|\mathcal{X}| = r$, $|\mathcal{Y}| = s$.

Assume without loss of generality that \mathcal{X} is universal. Let

$$f_1 \in \mathbb{B}^{\mathcal{X}} \text{ be such that } f_1([1 : r]) = (\underbrace{0, 0, \dots, 0}_{r-1 \text{ times}}, 1).$$

Such an f_1 exists because the set $[1 : r]$, as an index set of r consecutive integers, is a universal sampling set, so is in particular a sampling set for $\mathbb{B}^{\mathcal{X}}$. Similarly let

$$f_2 \in \mathbb{B}^{\mathcal{Y}} \text{ be such that } f_2([r : r + s - 1]) = (\underbrace{0, 0, \dots, 0}_{s-1 \text{ times}}, 1),$$

again possible because $[r : r + s - 1]$ is a set of s consecutive integers, and hence a sampling set for $\mathbb{B}^{\mathcal{Y}}$. Note that $f_1 f_2 \in \mathbb{B}^{\mathcal{X} + \mathcal{Y}}$ and so $|\mathcal{X} + \mathcal{Y}| \geq \text{supp}(\mathcal{F}(f_1 f_2))$. Note also that the zero set $\mathcal{Z}(f_1 f_2)$ of $f_1 f_2$ contains $[1 : r + s - 2]$, and hence, since the latter is a universal sampling set, $|\Omega(\mathcal{Z}(f_1 f_2))| \geq r + s - 2 = |\mathcal{X}| + |\mathcal{Y}| - 2$.

Now we apply the uncertainty principle of Theorem 9 to $f_1 f_2$. We have, so long as $f_1 f_2 \neq 0$,

$$\begin{aligned} |\mathcal{X} + \mathcal{Y}| &\geq \text{supp}(\mathcal{F}(f_1 f_2)) \\ &\geq 1 + |\Omega(\mathcal{Z}(f_1 f_2))| \\ &\geq 1 + |\mathcal{X}| + |\mathcal{Y}| - 2 = |\mathcal{X}| + |\mathcal{Y}| - 1, \end{aligned} \quad (52)$$

So we have proved that $|\mathcal{X} + \mathcal{Y}| \geq |\mathcal{X}| + |\mathcal{Y}| - 1$ if we know that $f_1 f_2 \neq 0$.

For this, again from Theorem 9 we have

$$|\mathcal{Z}(f_1)| \leq |\Phi(\text{supp}(\mathcal{F}f_1))| - 1 \leq |\Phi(\mathcal{X})| - 1,$$

since $f_1 \in \mathbb{B}^{\mathcal{X}}$. But \mathcal{X} is universal, so $\Phi(\mathcal{X}) = \mathcal{X}$ and

$$|\mathcal{Z}(f_1)| \leq |\mathcal{X}| - 1. \quad (53)$$

By definition of f_1 , the set $[1 : r-1] = [1 : |\mathcal{X}| - 1]$ is already in $\mathcal{Z}(f_1)$. Together with (53), this implies that f_1 cannot have any more zeros. In particular, $f_1(r+s-1) \neq 0$. Since $f_2(r+s-1) = 1$, $f_1 f_2$ cannot be identically zero and (52) applies. ■

An important generalization of the Cauchy-Davenport theorem to any finite abelian group, not necessarily of prime order, is due to Kneser, [18].

Theorem 13 (Kneser): Let G be a finite abelian group. Let $A, B \subseteq G$ be non empty subsets of G . Let H be the set of periods, defined by $H = \{h \in G : h + (A + B) = A + B\}$. (Thus $A + B$ is periodic if $H \neq \{0\}$.) Then

$$|A + B| \geq |A| + |B| - |H|.$$

Hence unless $A + B$ is periodic, $|A + B| \geq |A| + |B| - 1$.

Though the form is similar, this result neither implies nor is implied by Theorem 12. We give two examples. Let $N = 8$, $\mathcal{X} = \{0, 1\}$, $\mathcal{Y} = \{0, 4\}$. Then \mathcal{X} is universal and $\mathcal{X} + \mathcal{Y} = \{0, 1, 4, 5\}$ is periodic with period 4. So Theorem 12 applies, but Kneser's theorem does not. Next let $N = 16$, $\mathcal{X} = \{0, 2\}$, $\mathcal{Y} = \{0, 2, 4\}$. Then $\mathcal{X} + \mathcal{Y} = \{0, 2, 4, 6, 8, 10\}$, which is not periodic, and neither \mathcal{X} nor \mathcal{Y} is universal. So Kneser's theorem applies, but Theorem 12 does not. We hope to understand this more thoroughly.

APPENDIX A

CONDITION NUMBER ASSOCIATED WITH THE UNIVERSAL SAMPLING SET \mathcal{I}^*

An index set of consecutive integers is the simplest universal sampling set, but there is a catch in using it. Let \mathcal{I} be a universal sampling set of size d , $f \in \mathbb{C}^N$, and $f_{\mathcal{I}}$ the d -vector obtained from f by sampling at locations in \mathcal{I} . If f is in some bandlimited space $\mathbb{B}^{\mathcal{J}}$, $|\mathcal{J}| = d$, then the interpolation formula (4) reads

$$f = \mathcal{F}E_{\mathcal{J}}(E_{\mathcal{I}}^T \mathcal{F}E_{\mathcal{J}})^{-1} f_{\mathcal{I}}.$$

The practical difficulty is the computation of the inverse of $E_{\mathcal{I}}^T \mathcal{F}E_{\mathcal{J}}$. Suppose we use $\mathcal{I} = \mathcal{I}^* = [0 : d-1]$ as a universal sampling set. We give a lower bound on the condition number of $E_{\mathcal{I}}^T \mathcal{F}E_{\mathcal{J}}$ that can be quite large for some \mathcal{J} , even though the matrix $E_{\mathcal{I}}^T \mathcal{F}E_{\mathcal{J}}$ is invertible for all \mathcal{J} .

For $\mathcal{I} = [0 : d-1]$, note that

$$\begin{aligned} |\det(E_{\mathcal{I}}^T \mathcal{F}E_{\mathcal{J}})| &= |\det(\zeta_N^{ij})_{i \in \mathcal{I}, j \in \mathcal{J}}| \\ &= \prod_{j_1, j_2 \in \mathcal{J}} |\zeta_N^{j_1} - \zeta_N^{j_2}| \\ &= \prod_{j_1, j_2 \in \mathcal{J}} \left| 2 \sin \frac{2\pi(j_1 - j_2)}{N} \right|. \end{aligned}$$

If $\{\sigma_i\}$ are the singular values of $A = E_{\mathcal{I}}^T \mathcal{F}E_{\mathcal{J}}$, then

$$\det(A) = \sigma_1 \sigma_2 \sigma_3 \dots \sigma_d \geq \sigma_{\min}^d. \quad (54)$$

Also if $a_{rk} = \exp(-2\pi i r j_k / N)$ are the entries of A , then

$$d^2 = \sum_{r,k=0}^{d-1} |a_{rk}|^2 = \text{tr}(A^* A) = \sum_{r=0}^{d-1} \sigma_r^2 \leq d \sigma_{\max}^2, \quad (55)$$

and so $\sigma_{\max}^2 \geq d$.

From (54) and (55), the condition number satisfies

$$\frac{\sigma_{\max}}{\sigma_{\min}} \geq \sqrt{d} \left(\frac{1}{\prod_{j_1, j_2 \in \mathcal{J}} |2 \sin \frac{2\pi(j_1 - j_2)}{N}|} \right)^{1/2d}.$$

A possible scenario may be when d is very small and N is very large. In this case, the condition number can be very large if the frequency slots \mathcal{J} are clustered.

APPENDIX B

COUNTING BRACELETS

Several of our results, Theorem 4 for example, depend only on the bracelet of an index set rather than on the index set itself. Thus it is useful to know how many bracelets there are and how to enumerate them. Counting bracelets – actually, multicolored bracelets – is a standard application in combinatorics of the orbit stabilizer theorem, and the problem is treated in many places. Our situation is slightly different because we want a count that specifies the number of black beads in a black-and-white bracelet, corresponding to the size of the index set that determines the locations of the black beads. Nevertheless, the orbit stabilizer theorem can still be applied, and we have the following results.

Theorem 14: Let ϕ denote Euler's totient function. When N is odd, the number of black-and-white bracelets of length N with exactly d black beads is

$$\frac{1}{2} \binom{(N-1)/2}{d/2} + \frac{1}{2N} \sum_{k|N, k|d} \frac{\phi(k)}{N} \binom{N/k}{d/k} \quad \text{for even } d,$$

$$\frac{1}{2} \binom{(N-1)/2}{(d-1)/2} + \frac{1}{2N} \sum_{k|N, k|d} \frac{\phi(k)}{N} \binom{N/k}{d/k} \quad \text{for odd } d.$$

When N is even, the number of black-and-white bracelets of length N with exactly d black beads is

$$\frac{1}{2} \binom{N/2}{d/2} + \frac{1}{2N} \sum_{k|N, k|d} \frac{\phi(k)}{N} \binom{N/k}{d/k} \quad \text{for even } d,$$

$$\frac{1}{2} \binom{(N/2)-1}{(d-1)/2} + \frac{1}{2N} \sum_{k|N, k|d} \frac{\phi(k)}{N} \binom{N/k}{d/k} \quad \text{for odd } d.$$

We omit the proof; see [2]. An efficient algorithm for enumerating bracelets has been devised only recently by Sawada [19]. An algorithm for determining when two index sets are in the same necklace is due to J.P. Duval [20]. It can also be used for bracelets. See [2] for examples of both of these.

APPENDIX C

ADDITIONAL REFERENCES

Though our work has concerned discrete-time signals exclusively, there is also a notion of universal sampling sets for continuous-time signals. We will not give the definition; it is interesting and not clear what the relations between

the two may be. Here we cite only a few sources, starting with the paper of Landau [21] that featured the renowned necessary density condition on sampling sets. More recently, many interesting results have been obtained by Olevskii and Ulanovskii [22], [23] on universal sampling and stable reconstruction, by Matei and Meyer [24], who work with lattices and make contact with compressed sensing, and by Bass and Gröchenig [25], who consider random sampling. Of course, anyone writing on so fundamental a topic as sampling and interpolation will encounter an enormous literature, and most probably miss an equal or greater amount. We apologize to the authors of works we have missed.

ACKNOWLEDGMENTS

There are many people to thank for their interest, insight, and encouragement over quite some time, in particular S. Boyd, M. Chudnovsky, A. El Gamal, J.T. Gill, S. Gunturk, B. Hassibi, J. Sawada, J. Smith, and M. Tygert. We also thank the reviewers for their thorough and thoughtful comments.

REFERENCES

- [1] B. Osgood, A. Siripuram, and W. Wu, “Discrete sampling and interpolation: Orthogonal interpolating systems,” in preparation.
- [2] W. Wu, “Discrete sampling: Generalizations of the Nyquist-Shannon sampling theorem,” Ph.D. dissertation, Stanford University, 2010.
- [3] E. Candes, J. Romberg, and T. Tao, “Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on Information Theory*, vol. 52, (2), pp. 489–509, 2006.
- [4] R. Venkataramani and Y. Bresler, “Perfect reconstruction formulas and bounds on aliasing error in sub-nyquist nonuniform sampling of multiband signals,” *IEEE Transactions on Information Theory*, vol. 46, no. 6, pp. 2173–2183, 2000.
- [5] M. Mishali and Y. Eldar, “Blind multiband signal reconstruction: Compressed sensing for analog signals,” *IEEE Transactions on Signal Processing*, vol. 57, no. 3, pp. 993–1009, 2009.
- [6] T. Tao, “An uncertainty principle for cyclic groups of prime order,” *arXiv:math/0308286v6*.
- [7] A. Siripuram, “Sampling and interpolation of discrete signals: Orthogonality, universality and uncertainty,” Ph.D. dissertation, Stanford University.
- [8] V. V. Prasolov, *Problems and theorems in linear algebra*, ser. Translations of Mathematical Monographs. Providence, RI: Amer. Math. Soc., 1991, vol. 134.
- [9] P. Frenkel, “Simple proof of Chebotarev’s theorem on roots of unity,” *arXiv:math/0312398*.
- [10] S. Delvaux and M. Van Barel, “Rank-deficient submatrices of Fourier matrices,” *Linear Algebra Appl.*, vol. 429 (7), pp. 1587 – 1605, 2008.
- [11] O. Mitchell, “Note on determinants of powers,” *Amer. Jour. Math.*, vol. 4, no. 1, pp. 341–344, 1881.
- [12] R. Stanley, *Enumerative Combinatorics 2*, ser. Cambridge Studies in Advanced Mathematics. Cambridge, UK: Camb. Univ. Press., 1999, vol. 62.
- [13] R. Evans and I. Isaacs, “Generalized Vandermonde determinants and roots of unity of prime order,” *Proc. Amer. math. Soc.*, vol. 58, pp. 51–54, 1976.
- [14] D. Donoho and P. Stark, “Uncertainty principles and signal recovery,” *SIAM J. Appl. Math.*, vol. 49, no. 3, pp. 906–931, 1989.
- [15] C. Studer, P. Kuppinger, G. Pope, and H. Bolcskei, “Recovery of sparsely corrupted signals,” *arXiv.org/abs/1102.1621*.
- [16] B. Osgood, A. Siripuram, and W. Wu, “Additive uncertainty principles and signal reconstruction,” in preparation.
- [17] H. Davenport, “On the addition of residue classes,” *J. London Math. Soc.*, vol. 10, pp. 30–32, 1935.
- [18] M. Kneser, “Ein Satz über abelsche Gruppen mit Anwendungen auf die Geometrie der Zahlen,” *Math. Z.*, vol. 65, pp. 429–434, 1955.
- [19] J. Sawada, “Generating bracelets in amortized time,” *SIAM J. Comput.*, vol. 31, no. 1, pp. 259–268, 2001.
- [20] J. Duval, “Génération d’une section des classes de conjugaison et arbre des mots de Lyndon de longueur bornée,” *Theoretical Comp. Sci.*, vol. 60, no. 3, pp. 255–283, 1988.
- [21] H. Landau, “Necessary density conditions for sampling and interpolation of certain entire functions,” *Acta. Math.*, vol. 117, pp. 37–52, 1967.
- [22] A. Olevskii and A. Ulanovskii, “Universal sampling of band-limited signals,” *C.R. Math. Acad. Sci. Paris*, vol. 342, no. 12, pp. 927–931, 2006.
- [23] —, “Universal sampling and interpolation of bandlimited signals,” *Geom. funct. anal.*, vol. 18, pp. 1029–1052, 2008.
- [24] B. Matei and Y. Meyer, “A variant of compressed sensing,” *Rev. Mat. Iber.*, vol. 25, no. 2, pp. 669–692, 2009.
- [25] R. Bass and K. Gröchenig, “Random sampling of multivariate trigonometric polynomials,” *SIAM J. Math. Anal.*, vol. 36, p. 795, 2004.



Brad Osgood received his BS and MS at Carnegie-Mellon University and his Ph.D at the University of Michigan, all in mathematics. After a stint at Harvard he came to Stanford in 1985, first in the Mathematics Department and then in Electrical Engineering in the Information Systems Laboratory, where he is a Professor. Along with signal processing, his research in mathematics is in geometric function theory and differential geometry. Though becoming more digital, he plays trombone, the ultimate analog device.



Aditya Siripuram received his B.Tech and M.Tech degrees in Electrical Engineering from Indian Institute of Technology, Bombay in 2009. He is currently a PhD student in the Department of Electrical Engineering at Stanford University, and a recipient of the Stanford Graduate Fellowship. His interests include signal processing, coding theory and recreational mathematics.



William Wu received his B.Sc. in electrical engineering and computer science from the University of California, Berkeley, and his M.Sc. in electrical engineering, M.Sc. in mathematics, and Ph.D. in electrical engineering from Stanford University. His dissertation focused on sampling and reconstruction in finite dimensional signal spaces. Since 2010, he has been a member of the technical staff at the Jet Propulsion Laboratory (Pasadena, CA). His research interests include signal processing, information theory, scientific computation, and recreational math; he is the creator of wuriddles.com, an archive of puzzles.